arXiv:1703.00439v3 [math.OC] 16 Apr 2017

Doubly Accelerated Stochastic Variance Reduced Dual Averaging Method for Regularized Empirical Risk Minimization

Tomoya Murata [1] **Taiji Suzuki** [2]

Abstract

In this paper, we develop a new accelerated stochastic gradient method for efficiently solving the convex regularized empirical risk minimization problem in mini-batch settings. The use of mini-batches is becoming a golden standard in the machine learning community, because mini-batch settings stabilize the gradient estimate and can easily make good use of parallel computing. The core of our proposed method is the incorporation of our new "double acceleration" technique and variance reduction technique. We theoretically analyze our proposed method and show that our method much improves the mini-batch efficiencies of previous accelerated stochastic methods, and essentially only needs size $\sqrt{n}$ mini-batches for achieving the optimal iteration complexities for both non-strongly and strongly convex objectives, where n is the training set size. Further, we show that even in non-mini-batch settings, our method surpasses the best known convergence rate for non-strongly convex objectives, and it achieves the one for strongly convex objectives.

1. Introduction

We consider a composite convex optimization problem associated with regularized empirical risk minimization, which often arises in machine learning. In particular, our goal is to minimize the sum of finite smooth convex functions and a relatively simple (possibly) non-differentiable convex function by using first order methods in mini-batch settings. The use of mini-batches is becoming a golden standard in the machine learning community, because it is generally more efficient to execute matrix-vector multiplications over a mini-batch than an equivalent amount of vector-vector ones each over a single instance; and more importantly, mini-batch settings can easily make good use of parallel computing.

Traditional and effective methods for solving the above-mentioned problem are the "proximal gradient" (PG) method and "accelerated proximal gradient" (APG) method (Nesterov et al., 2007; Beck & Teboulle, 2009; Tseng, 2008). These methods are well known to achieve linear convergence for strongly convex objectives. Particularly, APG achieves optimal iteration complexities for both non-strongly and strongly convex objectives. However, these methods need a per iteration cost of $O(nd)$, where n denotes the number of components of the finite sum, and d is the dimension of the solution space. In typical machine learning tasks, n and d correspond to the number of instances and features respectively, which can be very large. Then, the per iteration cost of these methods can be considerably high.

A popular alternative is the "stochastic gradient descent" (SGD) method (Singer & Duchi, 2009; Hazan et al., 2007; Shalev-Shwartz & Singer, 2007). As the per iteration cost of SGD is only $O(d)$ in non-mini-batch settings, SGD is suitable for many machine learning tasks. However, SGD only achieves sublinear rates and is ultimately slower than PG and APG.

Recently, a number of stochastic gradient methods have been proposed; they use a variance reduction technique that utilizes the finite sum structure of the problem ("stochastic averaged gradient" (SAG) method (Roux et al., 2012; Schmidt et al., 2013), "stochastic variance reduced gradient" (SVRG) method (Johnson & Zhang, 2013; Xiao & Zhang, 2014) and SAGA (Defazio et al., 2014)). Even though the per iteration costs of these methods are same as that of SGD, they achieve a linear convergence for strongly convex objectives. Consequently, these methods dramatically improve the total computational cost of PG. However, in size b mini-batch settings, the rate is essentially b times worse than in non-mini-batch settings. This means that there is little benefit in applying mini-batch scheme to

[1]Department of Mathematical and Computing Sciences, Graduate School of Information Sciences and Engineering, Tokyo Institute of Technology [2]Department of Mathematical and Computing Science, School of Computing, Tokyo Institute of Technology. Correspondence to: Tomoya Murata <murata.t.ab@m.titech.ac.jp>, Taiji Suzuki <suzuki.t.ct@m.titech.ac.jp>.

these methods.

More recently, several authors have proposed accelerated stochastic methods for the composite finite sum problem ("accelerated stochastic dual coordinate ascent" (ASDCA) method (Shalev-Shwartz & Zhang, 2013), Universal Catalyst (UC) (Lin et al., 2015), "accelerated proximal coordinate gradient" (APCG) method (Lin et al., 2014a), "stochastic primal-dual coordinate" (SPDC) method (Zhang & Xiao, 2015), and Katyusha (Allen-Zhu, 2016)). ASDCA (UC), APCG, SPDC and Katyusha essentially achieve the optimal total computational cost[1] for strongly convex objectives[2] in non-mini-batch settings. However, in size b mini-batch settings, the rate is essentially $\sqrt{b}$ times worse than that in non-mini-batch settings, and these methods need size $O(n)$ mini-batches for achieving the optimal iteration complexity, which is essentially the same as APG. In addition, Nitanda (2014; 2015) has proposed the "accelerated mini-batch proximal stochastic variance reduced gradient" (AccProxSVRG) method and its variant, the "accelerated efficient mini-batch stochastic variance reduced gradient" (AMSVRG) method. In non-mini-batch settings, AccProxSVRG only achieves the same rate as SVRG. However, in mini-batch settings, AccProxSVRG significantly improves the mini-batch efficiency of non-accelerated variance reduction methods, and surprisingly, AccProxSVRG essentially only needs size $O(\sqrt{\kappa})$ mini-batches for achieving the optimal iteration complexity for strongly convex objectives, where κ is the condition number of the problem. However, the necessary size of mini-batches depends on the condition number and gradually increases when the condition number increases and ultimately matches with $O(n)$ for a large condition number.

Main contributions

We propose a new accelerated stochastic variance reduction method that achieves better convergence than existing methods do, and it particularly utilizes mini-batch settings well; it is called the "doubly accelerated stochastic variance reduced dual averaging" (DASVRDA) method. We describe the main features of our proposed method below and list the comparisons of our method with several preceding methods in Table 1.

- Our method significantly improves the mini-batch efficiencies of state-of-the-art methods. As a result, our method essentially only needs size $O(\sqrt{n})$ mini-batches[3] for achieving optimal iteration complexities for both non-strongly and strongly convex objectives.
- Even in non-mini-batch settings, our method improves the best-known convergence rate and is optimal up to a universal constant for non-strongly convex objectives, and it achieves the one for strongly convex objectives.

[1] More precisely, the rate of ASDCA (UC) is with extra log-factors, and near but worse than the one of APCG, SPDC and Katyusha. This means that ASDCA (UC) cannot be optimal.

[2] Katyusha also achieves a near optimal total computational cost for non-strongly convex objectives.

2. Preliminary

In this section, we provide several notations and definitions used in this paper. Then, we make assumptions for our analysis.

We use $\|\cdot\|$ to denote the Euclidean L_2 norm $\|\cdot\|_2$: $\|x\| = \|x\|_2 = \sqrt{\sum_i x_i^2}$. For natural number n, $[n]$ denotes set $\{1, \dots, n\}$.

Definition 2.1. We say that a function $f : \mathbb{R}^d \to \mathbb{R}$ is L-smooth ($L > 0$) if f is differentiable and satisfies

$$\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\| \ (\forall x, y \in \mathbb{R}^d). \tag{1}$$

If f is convex, (1) is equivalent to the following: (see Nesterov (2013)):

$$f(x) + \langle \nabla f(x), y - x\rangle + \frac{1}{2L}\|\nabla f(x) - \nabla f(y)\|^2 \leq f(y) \tag{2}$$

$(\forall x, y \in \mathbb{R}^d)$.

Definition 2.2. A convex function $f : \mathbb{R}^d \to \mathbb{R}$ is called μ-strongly convex ($\mu \geq 0$) if f satisfies

$$\frac{\mu}{2}\|x - y\|^2 + \langle \xi, y - x\rangle + f(x) \leq f(y),$$
$$(\forall x, y \in \mathbb{R}^d, \forall \xi \in \partial f(x))$$

where $\partial f(x)$ denotes the set of the subgradients of f at x.

Note that if f is μ-strongly convex, then f has the unique minimizer.

Definition 2.3. We say that a function $f : \mathbb{R}^d \to \mathbb{R}$ is μ-optimally strongly convex ($\mu \geq 0$) if f has a minimizer and satisfies

$$\frac{\mu}{2}\|x - x_*\|^2 \leq f(x) - f(x_*)$$
$$(\forall x \in \mathbb{R}^d, \forall x_* \in \operatorname{argmin}_{x \in \mathbb{R}^d} f(x)).$$

[3] Actually, when $L/\varepsilon \leq n$ and $L/\mu \leq n$, our method needs size $O(n\sqrt{\varepsilon/L})$ and $O(n\sqrt{\mu/L})$ mini-batches, respectively, which are larger than $O(\sqrt{n})$, but smaller than $O(n)$. Achieving optimal iteration complexity for solving high accuracy and bad conditioned problems is much more important than doing ones with low accuracy and well-conditioned ones, because the former needs more overall computational cost than the latter.

	μ-strongly convex			Non-strongly convex		
	Total computational cost in size b mini-batch settings	Necessary size of mini-batches When $L/\mu \geq n$	Necessary size of mini-batches $L/\mu \leq n$	Total computational cost in size b mini-batch settings	Necessary size of mini-batches When $L/\varepsilon \geq n$	Necessary size of mini-batches $L/\varepsilon \leq n$
SVRG (SVRG^{++})	$O\left(d\left(n+\frac{bL}{\mu}\right)\log\left(\frac{1}{\varepsilon}\right)\right)$	Unattainable	Unattainable	$O\left(d\left(n\log\left(\frac{1}{\varepsilon}\right)+\frac{bL}{\varepsilon}\right)\right)$	Unattainable	Unattainable
ASDCA (UC)	$\widetilde{O}\left(d\left(n+\sqrt{\frac{nbL}{\mu}}\right)\log\left(\frac{1}{\varepsilon}\right)\right)$	Unattainable	Unattainable	$\widetilde{O}\left(d\left(\frac{n+\sqrt{nbL}}{\sqrt{\varepsilon}}\right)\right)$	Unattainable	Unattainable
APCG	$O\left(d\left(n+\sqrt{\frac{nbL}{\mu}}\right)\log\left(\frac{1}{\varepsilon}\right)\right)$	$O(n)$	$O(n)$	No direct analysis	Unattainable	Unattainable
SPDC	$O\left(d\left(n+\sqrt{\frac{nbL}{\mu}}\right)\log\left(\frac{1}{\varepsilon}\right)\right)$	$O(n)$	$O(n)$	No direct analysis	Unattainable	Unattainable
Katyusha	$O\left(d\left(n+\sqrt{\frac{nbL}{\mu}}\right)\log\left(\frac{1}{\varepsilon}\right)\right)$	$O(n)$	$O(n)$	$O\left(d\left(n\log\left(\frac{1}{\varepsilon}\right)+\sqrt{\frac{nbL}{\varepsilon}}\right)\right)$	$O(n)$	$O(n)$
AccProxSVRG	$O\left(d\left(n+\left(\frac{n-b}{n-1}\right)\frac{L}{\mu}+b\sqrt{\frac{L}{\mu}}\right)\log\left(\frac{1}{\varepsilon}\right)\right)$	$O\left(\sqrt{\frac{L}{\mu}}\right)$	$O\left(n\sqrt{\frac{\mu}{L}}\right)$	No direct analysis	Unattainable	Unattainable
DASVRDA	$O\left(d\left(n+\sqrt{\frac{nL}{\mu}}+b\sqrt{\frac{L}{\mu}}\right)\log\left(\frac{1}{\varepsilon}\right)\right)$	$O\left(\sqrt{n}\right)$	$O\left(n\sqrt{\frac{\mu}{L}}\right)$	$O\left(d\left(n+\sqrt{\frac{nL}{\varepsilon}}+b\sqrt{\frac{L}{\varepsilon}}\right)\right)$	$O\left(\sqrt{n}\right)$	$O\left(n\sqrt{\frac{\varepsilon}{L}}\right)$

Table 1. Comparisons of our method with SVRG (SVRG^{++} (Allen-Zhu & Yuan, 2016)), ASDCA (UC), APCG, SPDC, Katyusha and AccProxSVRG. n is the number of components of the finite sum, d is the dimension of the solution space, b is the mini-batch size, L is the smoothness parameter of the finite sum, μ is the strong convexity parameter of objectives (see Def. 2.1 and Def. 2.2 for their definitions), and ε is accuracy. "Necessary size of mini-batches " indicates the order of the necessary size of mini-batches for achieving optimal iteration complexities $O(\sqrt{L/\mu}\log(1/\varepsilon))$ and $O(\sqrt{L/\varepsilon})$ for strongly and non-strongly convex objectives, respectively. We regard one computation of a full gradient as n/b iterations in size b mini-batch settings, for a fair comparison. "Unattainable" implies that the algorithm cannot achieve the optimal iteration complexity even if it uses size n mini-batches. $\widetilde{O}$ hides extra log-factors. The results marked in red denote the main contributions of this paper.

If a function f is μ-strongly convex, then f is clearly μ-optimally strongly convex.

Definition 2.4. We say that a convex function $f:\mathbb{R}^d\to\mathbb{R}$ is relatively simple if computing the proximal mapping of f at y,

$$\mathrm{prox}_f(y)=\operatorname*{argmin}_{x\in\mathbb{R}^d}\left\{\frac{1}{2}\|x-y\|^2+f(x)\right\},$$

takes at most $O(d)$ computational cost, for any $y\in\mathbb{R}^d$.

As f is convex, function $(1/2)\|x-y\|^2+f(x)$ is 1-strongly convex, and function prox_f is well-defined.

Now, we formally describe the problem to be considered in this paper and the assumptions for our theory. In this paper, we consider the following composite convex minimization problem:

$$\min_{x\in\mathbb{R}^d}\ \{P(x)\stackrel{\mathrm{def}}{=}F(x)+R(x)\},\tag{3}$$

where $F(x)=\frac{1}{n}\sum_{i=1}^n f_i(x)$. Here each $f_i:\mathbb{R}^d\to\mathbb{R}$ is a L_i-smooth convex function and $R:\mathbb{R}^d\to\mathbb{R}$ is a relatively simple and (possibly) non-differentiable convex function. Problems of this form often arise in machine learning and fall under regularized empirical risk minimization (ERM). In ERM problems, we are given n training examples $\{(a_i,b_i)\}_{i=1}^n$, where each $a_i\in\mathbb{R}^d$ is the feature vector of example i, and each $b_i\in\mathbb{R}$ is the label of example i. The following regression and classification problems are typical examples of ERM on our setting:

- Lasso: $f_i(x)=\frac{1}{2}(a_i^\top x-b_i)^2$ and $R(x)=\lambda\|x\|_1$.
- Ridge logistic regression: $f_i(x)=\log(1+\exp(-b_ia_i^\top x))$ and $R(x)=\frac{\lambda}{2}\|x\|_2^2$.
- Support vector machines: $f_i(x)=\bar{h}_i^\nu(a_i^\top x)$ and $R(x)=\frac{\lambda}{2}\|x\|_2^2$.

Here, $\bar{h}_i^\nu$ is a smooth variant of hinge loss (for the definition, for example, see Shalev-Shwartz & Zhang (2013)).

We make the following assumptions for our analysis:

Assumption 1. *There exists a minimizer x_* of (3).*

Assumption 2. *Each f_i is convex and L_i-smooth.*

The above examples satisfy this assumption with $L_i=\|a_i\|_2^2$ for the squared loss, $L_i=\|a_i\|_2^2/4$ for the logistic loss, and $L_i=\nu\|a_i\|_2^2$ for the smoothed hinge loss.

Assumption 3. *Regularization function R is convex and is relatively simple.*

For example, Elastic Net regularizer $R(x)=\lambda_1\|x\|_1+(\lambda_2/2)\|x\|_2^2$ $(\lambda_1,\lambda_2\geq 0)$ satisfies this assumption. Indeed, we can analytically compute the proximal mapping of R by $\mathrm{prox}_R(z)=((1/(1+\lambda_2))\mathrm{sign}(z_j)\max\{|z_j|-\lambda_1,0\})_{j=1}^d$, and this costs only $O(d)$.

We always consider Assumption 1, 2 and 3 in this paper.

Assumption 4. *There exists $\mu>0$ such that objective function P is μ-optimally strongly convex.*

If each f_i is convex, then for Elastic Net regularization function $R(x)=\lambda_1\|x\|_1+(\lambda_2/2)\|x\|_2^2$ ($\lambda_1\geq 0$ and $\lambda_2>0$), Assumption 4 holds with $\mu=\lambda_2$.

We further consider Assumption 4 when we deal with strongly convex objectives.

Algorithm 1 PG $(\widetilde{x}_0, \eta, S)$

Input: $\widetilde{x}_0 \in \mathbb{R}^d$, $\eta > 0$, $S \in \mathbb{N}$.
  for $s = 1$ to S **do**
    $\widetilde{x}_s = \text{One Stage PG}(\widetilde{x}_{s-1}, \eta)$.
  end for
Output: $\frac{1}{S}\sum_{s=1}^{S} \widetilde{x}_s$.

Algorithm 2 One Stage PG $(\widetilde{x}, \eta)$

Input: $\widetilde{x} \in \mathbb{R}^d$, $\eta > 0$.
  $\widetilde{x}^+ = \text{prox}_{\eta R}(\widetilde{x} - \eta \nabla F(\widetilde{x}))$.
Output: $\widetilde{x}^+$.

3. Our Approach: Double Acceleration

In this section, we provide high-level ideas of our main contribution called "double acceleration."

First, we consider deterministic PG (Algorithm 1) and (non-mini-batch) SVRG (Algorithm 3). PG is an extension of the steepest descent to proximal settings. SVRG is a stochastic gradient method using the variance reduction technique, which utilizes the finite sum structure of the problem, and it achieves a faster convergence rate than PG does. As SVRG (Algorithm 3) matches with PG (Algorithm 1) when the number of inner iterations is $m = 1$, SVRG can be seen as a generalization of PG. The key element of SVRG is employing a simple but powerful technique called the variance reduction technique for gradient estimate. The variance reduction of the gradient is realized by setting $g_k = \nabla f_{i_k}(x_{k-1}) - \nabla f_{i_k}(\widetilde{x}) + \nabla F(\widetilde{x})$ rather than vanilla stochastic gradient $\nabla f_{i_k}(x_{k-1})$. Generally, stochastic gradient $\nabla f_{i_k}(x_{k-1})$ is an unbiased estimator of $\nabla F(x_{k-1})$, but it may have high variance. In contrast, g_k is also unbiased, and one can show that its variance is "reduced"; that is, the variance converges to zero as x_{k-1} and $\widetilde{x}$ to x_*.

Next, we explain to the method of accelerating SVRG and obtaining an even faster convergence rate based on our new but quite natural idea "outer acceleration." First, we would like to remind you that the procedure of deterministic APG is given as described in Algorithm 5. APG uses the famous "momentum" scheme and achieves the optimal iteration complexity. Our natural idea is replacing One Stage PG in Algorithm 5 with One Stage SVRG. With slight modifications, we can show that this algorithm improves the rates of PG, SVRG and APG, and is optimal. We call this new algorithm outerly accelerated SVRG (Note that the algorithm matches with APG when $m = 1$ and thus, can be seen as a direct generalization of APG). However, this algorithm has poor mini-batch efficiency, because in size b mini-batch settings, the rate of this algorithm is essentially $\sqrt{b}$ times worse than that of non-mini-batch settings. State-of-the-art methods APCG, SPDC, and Katyusha also suffer from the same problem in the mini-batch setting.

Algorithm 3 SVRG $(\widetilde{x}_0, \eta, m, S)$

Input: $\widetilde{x}_0 \in \mathbb{R}^d$, $\eta > 0$, $m, S \in \mathbb{N}$.
  for $s = 1$ to S **do**
    $\widetilde{x}_s = \text{One Stage SVRG}(\widetilde{x}_{s-1}, \eta, m)$.
  end for
Output: $\frac{1}{S}\sum_{s=1}^{S} \widetilde{x}_s$.

Algorithm 4 One Stage SVRG $(\widetilde{x}, \eta, m)$

Input: $\widetilde{x} \in \mathbb{R}^d$, $\eta > 0$, $m \in \mathbb{N}$.
  $x_0 = \widetilde{x}$.
  for $k = 1$ to m **do**
    Pick $i_k \in \{1, \ldots, n\}$ randomly.
    $g_k = \nabla f_{i_k}(x_{k-1}) - \nabla f_{i_k}(\widetilde{x}) + \nabla F(\widetilde{x})$.
    $x_k = \text{prox}_{\eta R}(x_{k-1} - \eta g_k)$.
  end for
Output: $\frac{1}{n}\sum_{k=1}^{n} x_k$.

Now, we illustrate that for improving the mini-batch efficiency, using the "inner acceleration" technique is beneficial. Nitanda (Nitanda, 2014) has proposed AccProxSVRG in mini-batch settings. He applied the momentum scheme to One Stage SVRG, and we call this technique "inner" acceleration. He showed that the inner acceleration could significantly improve the mini-batch efficiency of vanilla SVRG. This fact indicates that inner acceleration is essential to fully utilize the mini-batch settings. However, AccProxSVRG is not a truly accelerated method, because in non-mini-batch settings, the rate of AccProxSVRG is same as that of vanilla SVRG.

In this way, we arrive at our main high-level idea called "double" acceleration, which involves applying momentum scheme to both outer and inner algorithms. This enables us not only to lead to the optimal total computational cost in non-mini-batch settings, but also to improving the mini-batch efficiency of vanilla acceleration methods.

We have considered SVRG and its accelerations so far; however, we actually adopt stochastic variance reduced dual averaging (SVRDA) rather than SVRG itself, because we can construct lazy update rules of (innerly) accelerated SVRDA for sparse data.

4. Algorithm Description

In this section, we describe the concrete procedure of the proposed algorithm in detail.

We provide details of the doubly accelerated SVRDA (DASVRDA) method for non-strongly convex objectives in Algorithm 6. Our momentum step is slightly different from that of vanilla deterministic accelerated methods: we not only add momentum term $((\widetilde{\theta}_{s-1} - 1)/\widetilde{\theta}_s)(\widetilde{x}_{s-1} -$

Algorithm 5 APG $(\widetilde{x}_0, \eta, S)$

Input: $\widetilde{x}_0 \in \mathbb{R}^d$, $\eta > 0$, $S \in \mathbb{N}$.
  $\widetilde{x}_{-1} = \widetilde{x}_0$.
  $\widetilde{\theta}_0 = 0$.
  for $s = 1$ to S **do**
    $\widetilde{\theta}_s = \frac{s+1}{2}$.
    $\widetilde{y}_s = \widetilde{x}_{s-1} + \frac{\widetilde{\theta}_{s-1}-1}{\widetilde{\theta}_s}(\widetilde{x}_{s-1} - \widetilde{x}_{s-2})$.
    $\widetilde{x}_s = $ One Stage PG$(\widetilde{y}_s, \eta)$.
  end for
Output: x_S.

$\widetilde{x}_{s-2}$) to the current solution $\widetilde{x}_{s-1}$ but also add term $(\widetilde{\theta}_{s-1}/\widetilde{\theta}_s)(\widetilde{z}_{s-1} - \widetilde{x}_{s-1})$, where $\widetilde{z}_{s-1}$ is the current more "aggressively" updated solution rather than $\widetilde{x}_{s-1}$; thus, this term also can be interpreted as momentum[4]. Then, we feed $\widetilde{y}_s$ to One Stage Accelerated SVRDA (Algorithm 7) as an initial point. Note that Algorithm 6 can be seen as a direct generalization of APG, because if we set $m = 1$, One Stage Accelerated SVRDA is essentially the same as one iteration PG with initial point $\widetilde{y}_s$; then, we can see that $\widetilde{z}_s = \widetilde{x}_s$, and Algorithm 6 essentially matches with deterministic APG. Next, we move to One Stage Accelerated SVRDA (Algorithm 7). Algorithm 7 is essentially a combination of the "accelerated regularized dual averaging" (AccSDA) method (Xiao, 2009) with the variance reduction technique of SVRG. It updates z_k by using the weighted average of all past variance reduced gradients $\bar{g}_k$ instead of only using a single variance reduced gradient g_k. Note that for constructing variance reduced gradient g_k, we use the full gradient of F at $\widetilde{x}_{s-1}$ rather than the initial point $\widetilde{y}_s$.

Remark. In Algorithm 7, we pick b indexes according to i.i.d. non-uniform distribution $Q = \{q_i\} = \left\{\frac{L_i}{n\bar{L}}\right\}$. Instead of that, we can pick b indexes, such that each index i_k^ℓ is uniformly picked from B^ℓ, where $\{B^\ell\}_{\ell=1}^b$ is the predefined disjoint partition of $[n]$ with size $|B^\ell| = n/b$. If we adopt this scheme, when we parallelize the algorithm using b machines, each machine only needs to store the corresponding partition of the data set rather than the whole dataset, and this can reduce communication cost and memory cost. In this setting, the convergence analysis in Section 5 can be easily revised by simply replacing $\bar{L}$ with $L_{\max} = \max_{i\in[n]} L_i$.

Algorithm 8 is our proposed method for strongly convex objectives. Instead of directly accelerating the algorithm using a constant momentum rate, we restart Algorithm 6. Restarting scheme has several advantages both theoretically and practically. First, the restarting scheme only requires the optimal strong convexity of the objective (Def. 2.3) instead of the ordinary strong convexity (Def. 2.2). Whereas, non-restarting accelerated algorithms essentially require the ordinary strong convexity of the objective. Second, for restarting algorithms, we can utilize adaptive restart schemes (O'Donoghue & Candes, 2015). The adaptive restart schemes have been originally proposed for deterministic cases. The schemes are heuristic but quite effective empirically. The most fascinating property of these schemes is that we need not prespecify the strong convexity parameter μ, and the algorithms adaptively determine the restart timings. O'Donoghue & Candes (2015) have proposed two heuristic adaptive restart schemes: the function scheme and gradient scheme. We can easily apply these ideas to our method, because our method is a direct generalization of the deterministic APG. For the function scheme, we restart Algorithm 6 if $P(\widetilde{x}_s) > P(\widetilde{x}_{s-1})$. For the gradient scheme, we restart the algorithm if $(\widetilde{y}_s - \widetilde{x}_s)^\top(\widetilde{y}_{s+1} - \widetilde{x}_s) > 0$. Here $\widetilde{y}_s - \widetilde{x}_s$ can be interpreted as a "one stage" gradient mapping of P at $\widetilde{y}_s$. As $\widetilde{y}_{s+1} - \widetilde{x}_s$ is the momentum, this scheme can be interpreted such that we restart whenever the momentum and negative one Stage gradient mapping form an obtuse angle (this means that the momentum direction seems to be "bad"). We numerically demonstrate the effectiveness of these schemes in section 7.

Parameter tunings

For DASVRDA$^{\text{ns}}$, only learning rate η needs to be tuned, because we can theoretically obtain the optimal choice of γ, and we can naturally use $m = n/b$ as a default epoch length (see Section 5). For DASVRDA$^{\text{sc}}$, both learning rate η and fixed restart interval S need to be tuned.

5. Convergence Analysis of DASVRDA Method

In this section, we provide the convergence analysis of our algorithms. Unless otherwise specified, serial computation is assumed. First, we consider the DASVRDA$^{\text{ns}}$ algorithm.

Theorem 5.1. *Suppose that Assumptions 1, 2 and 3 hold. Let $\widetilde{x}_0 \in \mathbb{R}^d$, $\gamma > 1$, $m \in \mathbb{N}$, $b \in [n]$ and $S \in \mathbb{N}$. Then DASVRDA$^{\text{ns}}(\widetilde{x}_0, \gamma, \{L_i\}_{i=1}^n, m, b, S)$ satisfies*

$$\mathbb{E}\left[P(\widetilde{x}_S) - P(x_*)\right] \leq \frac{\left(1 + \frac{\gamma(m+1)}{b}\right)\bar{L}}{\left(1 - \frac{1}{\gamma}\right)^2 (S+1)^2 m^2}\|\widetilde{x}_0 - x_*\|^2.$$

The proof of Theorem 5.1 is found in the supplementary

[4] This form also arises in Monotone APG (Li & Lin, 2015). In Algorithm 7, $\widetilde{x} = x_m$ can be rewritten as $(2/(m(m+1)))\sum_{k=1}^m k z_k$, which is a weighted average of z_k; thus, we can say that $\widetilde{z}$ is updated more "aggressively" than $\widetilde{x}$. For the outerly accelerated SVRG (that is a combination of Algorithm 6 with vanilla SVRG, see section 3), $\widetilde{z}$ and $\widetilde{x}$ correspond to x_m and $(1/m)\sum_{k=1}^m x_k$ in (Xiao & Zhang, 2014), respectively. Thus, we can also see that $\widetilde{z}$ is updated more "aggressively" than $\widetilde{x}$.

Algorithm 6 DASVRDA$^{\mathrm{ns}}(\widetilde{x}_0, \gamma, \{L_i\}_{i=1}^n, m, b, S)$

Input: $\widetilde{x}_0 \in \mathbb{R}^d$, $\gamma > 1$, $\{L_i > 0\}_{i=1}^n$, $m \in \mathbb{N}$, $S \in \mathbb{N}$, $b \in [n]$.
$\widetilde{x}_{-1} = \widetilde{z}_0 = \widetilde{x}_0$.
$\widetilde{\theta}_0 = 0$.
$\bar{L} = \frac{1}{n}\sum_{i=1}^n L_i$.
$Q = \{q_i\} = \left\{\frac{L_i}{n\bar{L}}\right\}$.
$\eta = \frac{1}{\left(1+\frac{\gamma(m+1)}{b}\right)\bar{L}}$.
for $s = 1$ to S **do**
$\widetilde{\theta}_s = \left(1 - \frac{1}{\gamma}\right)\frac{s+1}{2}$.
$\widetilde{y}_s = \widetilde{x}_{s-1} + \frac{\widetilde{\theta}_{s-1}-1}{\widetilde{\theta}_s}(\widetilde{x}_{s-1} - \widetilde{x}_{s-2}) + \frac{\widetilde{\theta}_{s-1}}{\widetilde{\theta}_s}(\widetilde{z}_{s-1} - \widetilde{x}_{s-1})$.
$(\widetilde{x}_s, \widetilde{z}_s) =$ One Stage AccSVRDA$(\widetilde{y}_s, \widetilde{x}_{s-1}, \eta, m, b, Q)$.
end for
Output: $\widetilde{x}_S$.

material. We can easily see that the optimal choice of γ is $(3+\sqrt{9+8b/(m+1)})/2 = O(1+b/m)$ (see supplementary material). We denote this value as γ_*. From Theorem 5.1, we obtain the following corollary:

Corollary 5.2. *Suppose that Assumptions 1, 2, and 3 hold. Let $\widetilde{x}_0 \in \mathbb{R}^d$, $\gamma = \gamma_*$, $m \in \mathbb{N}$, $b \in [n]$ and $S \in \mathbb{N}$. Then, DASVRDA$^{\mathrm{ns}}(\widetilde{x}_0, \gamma_*, \{L_i\}_{i=1}^n, m, b, S)$ satisfies*

$$\mathbb{E}\left[P(\widetilde{x}_S) - P(x_*)\right] \leq O\left(\frac{\left(1+\frac{m}{b}\right)\bar{L}}{S^2 m^2}\|\widetilde{x}_0 - x_*\|^2\right).$$

In other words, to achieve $\mathbb{E}\left[P(\widetilde{x}_S) - P(x_*)\right] \leq \varepsilon$, it suffices to execute $S = O(1 + (1/m + 1/\sqrt{mb})\sqrt{\bar{L}\|\widetilde{x}_0 - x_*\|^2/\varepsilon})$ iterations, and thus, DASVRDA$^{\mathrm{ns}}(\widetilde{x}_0, \gamma_*, \{L_i\}_{i=1}^n, m, b, S)$ achieves a total computational cost of

$$O\left(d\left((n+mb) + (n+mb)\left(\frac{1}{m} + \frac{1}{\sqrt{mb}}\right)\sqrt{\frac{\bar{L}\|\widetilde{x}_0 - x_*\|^2}{\varepsilon}}\right)\right)$$

Furthermore, if $m \propto n/b$, we obtain

$$O\left(d\left(n + \left(b + \sqrt{n}\right)\sqrt{\frac{\bar{L}\|\widetilde{x}_0 - x_*\|^2}{\varepsilon}}\right)\right).$$

Remark. Corollary 5.2 implies that if the mini-batch size b is $O(\sqrt{n})$, DASVRDA$^{\mathrm{ns}}(\widetilde{x}_0, \gamma_*, \{L_i\}_{i=1}^n, n/b, b, S)$ still achieves the total computational cost of $O(d(n+\sqrt{n\bar{L}/\varepsilon}))$, which is better than $O(d(n\log(1/\varepsilon) + \sqrt{nb\bar{L}/\varepsilon}))$ of Katyusha even if $b = 1$. The rate is optimal up to a universal constant, when compared to the lower bound of Theorem 7 in (Woodworth & Srebro, 2016).

Remark. Corollary 5.2 also implies that DASVRDA$^{\mathrm{ns}}$ only needs size $O(\sqrt{n})$ mini-batches for achieving the optimal iteration complexity of $O(\sqrt{L/\varepsilon})$, when $L/\varepsilon \geq n$. In contrast, Katyusha needs size $O(n)$ mini-batches for achieving the optimal iteration complexity. Note that even when $L/\varepsilon \leq n$, our method only needs size $O(n\sqrt{\varepsilon/L})$ mini-batches[5], that is smaller than $O(n)$ of Katyusha.

Algorithm 7 One Stage AccSVRDA $(\widetilde{y}, \widetilde{x}, \eta, m, b, Q)$

Input: $\widetilde{y}, \widetilde{x}, \eta > 0$, $m \in \mathbb{N}$, $b \in [n]$, Q.
$x_0 = z_0 = \widetilde{y}$.
$\bar{g}_0 = 0$.
$\theta_0 = \frac{1}{2}$.
for $k = 1$ to m **do**
Pick independently $i_k^1, \ldots, i_k^b \sim Q$, $I_k = \{i_k^\ell\}_{\ell=1}^b$.
$\theta_k = \frac{k+1}{2}$.
$y_k = \left(1 - \frac{1}{\theta_k}\right)x_{k-1} + \frac{1}{\theta_k}z_{k-1}$.
$g_k = \frac{1}{b}\sum_{i\in I_k}\frac{1}{nq_i}\left(\nabla f_i(y_k) - \nabla f_i(\widetilde{x})\right) + \nabla F(\widetilde{x})$.
$\bar{g}_k = \left(1 - \frac{1}{\theta_k}\right)\bar{g}_{k-1} + \frac{1}{\theta_k}g_k$.
$z_k = \operatorname*{argmin}_{z\in\mathbb{R}^d}\left\{\langle \bar{g}_k, z\rangle + R(z) + \frac{1}{2\eta\theta_k\theta_{k-1}}\|z - z_0\|^2\right\}$
$= \mathrm{prox}_{\eta\theta_k\theta_{k-1}R}\left(z_0 - \eta\theta_k\theta_{k-1}\bar{g}_k\right)$.
$x_k = \left(1 - \frac{1}{\theta_k}\right)x_{k-1} + \frac{1}{\theta_k}z_k$.
end for
Output: (x_m, z_m).

Algorithm 8 DASVRDA$^{\mathrm{sc}}(\check{x}_0, \gamma, \{L_i\}_{i=1}^n, m, b, S, T)$

Input: $\check{x}_0 \in \mathbb{R}^d$, $\gamma \geq 1$, $\{L_i > 0\}_{i=1}^n$, $m \in \mathbb{N}$, $b \in [n]$, $S, T \in \mathbb{N}$.
for $t = 1$ to T **do**
$\check{x}_t =$ DASVRDA$^{\mathrm{ns}}(\check{x}_{t-1}, \gamma, \{L_i\}_{i=1}^n, m, b, S)$.
end for
Output: $\check{x}_T$.

Next, we analyze the DASVRDA$^{\mathrm{sc}}$ algorithm for optimally strongly convex objectives. Combining Theorem 5.1 with the optimal strong convexity of the objective function immediately yields the following theorem, which implies that the DASVRDA$^{\mathrm{sc}}$ algorithm achieves a linear convergence.

Theorem 5.3. *Suppose that Assumptions 1, 2, 3 and 4 hold. Let $\check{x}_0 \in \mathbb{R}^d$, $\gamma = \gamma_*$, $m \in \mathbb{N}$, $b \in [n]$ and $T \in \mathbb{N}$. Define $\rho \stackrel{\mathrm{def}}{=} \{2(1+\gamma_*(m+1)/b)\bar{L}\}/\{(1-1/\gamma_*)^2(S+1)^2m^2\mu\}$. If S is sufficiently large such that $\rho \in (0,1)$, then DASVRDA$^{\mathrm{sc}}(\check{x}_0, \gamma_*, \{L_i\}_{i=1}^n, m, b, S, T)$ satisfies*

$$\mathbb{E}[P(\check{x}_T) - P(x_*)] \leq \rho^T[P(\check{x}_0) - P(x_*)].$$

From Theorem 5.3, we have the following corollary.

Corollary 5.4. *Suppose that Assumptions 1, 2, 3 and 4 hold. Let $\check{x}_0 \in \mathbb{R}^d$, $\gamma = \gamma_*$, $m \propto n/b$, $b \in [n]$. There exists $S = O(1 + (b/n + 1/\sqrt{n})\sqrt{\bar{L}/\mu})$, such that $1/\log(1/\rho) = O(1)$. Moreover, if we appropriately choose $T = O(\log(P(\check{x}_0) - P(x_*)/\varepsilon)$, then a total computa-*

[5]Note that we regard one computation of a full gradient as n/b iterations in size b mini-batch settings.

tional cost of DASVRDA$^{\mathrm{sc}}(\check{x}_0, \gamma_*, \{L_i\}_{i=1}^n, m, b, S, T)$ *for* $\mathbb{E}\left[P(\check{x}_T) - P(x_*)\right] \leq \varepsilon$ *is*

$$O\left(d\left(n + \left(b + \sqrt{n}\right)\sqrt{\frac{\bar{L}}{\mu}}\right)\log\left(\frac{P(\check{x}_0) - P(x_*)}{\varepsilon}\right)\right).$$

Remark. Corollary 5.4 implies that if the mini-batch size b is $O(\sqrt{n})$, DASVRDA$^{\mathrm{sc}}(\check{x}_0, \gamma_*, \{L_i\}_{i=1}^n, n/b, b, S, T)$ still achieves the total computational cost of $O(d(n + \sqrt{n\bar{L}/\mu})\log(1/\varepsilon))$, which is much better than $O(d(n + \sqrt{nb\bar{L}/\mu})\log(1/\varepsilon))$ of APCG, SPDC, and Katyusha.

Remark. Corollary 5.4 also implies that DASVRDA$^{\mathrm{sc}}$ only needs size $O(\sqrt{n})$ mini-batches for achieving the optimal iteration complexity $O(\sqrt{L/\mu}\log(1/\varepsilon))$, when $L/\mu \geq n$. In contrast, APCG, SPDC and Katyusha need size $O(n)$ mini-batches and AccProxSVRG does $O(\sqrt{L/\mu})$ ones for achieving the optimal iteration complexity. Note that even when $L/\mu \leq n$, our method only needs size $O(n\sqrt{\mu/L})$ mini-batches [6]. This size is smaller than $O(n)$ of APCG, SPDC, and Katyusha, and the same as that of AccProxSVRG.

6. Efficient Implementation for Sparse Data: Lazy Update

In this section, we briefly comment on the sparse implementations of our algorithms.

Originally, lazy update was proposed in online settings (Duchi et al., 2011). Generally, it is difficult for accelerated stochastic variance reduction methods to construct lazy update rules because (i) generally, variance reduced gradients are not sparse even if stochastic gradients are sparse; (ii) if we adopt the momentum scheme, the updated solution becomes a convex combination of previous solutions; and (iii) for non-strongly convex objectives, the momentum rate must not be constant. Konečnỳ et al. (2016) have tackled the problem of (i) on non-accelerated settings and derived lazy update rules of the "mini-batch semi-stochastic gradient descent" (mS2GD) method. Allen-Zhu (2016) has only mentioned that lazy updates can be applied to Katyusha but did not give explicit lazy update rules of Katyusha. Particularly, for non-strongly convex objectives, it seems to be difficult to derive lazy update rules owing to the difficulty of (iii). The reason we adopt the stochastic dual averaging scheme (Xiao, 2009) rather than stochastic gradient descent for our method is to be able to overcome the difficulties faced in (i), (ii), and (iii). The lazy update rules of our method support both non-strongly and strongly convex objectives. The formal lazy update algorithms of our method can be found in the supplementary material.

[6] Note that the required size is $O(n\sqrt{\mu/L})(\leq O(n))$, which is not $O(n\sqrt{L/\mu}) \geq O(n)$.

7. Numerical Experiments

In this section, we provide numerical experiments to demonstrate the performance of DASVRDA.

We numerically compare our method with several well-known stochastic gradient methods in mini-batch settings: SVRG (Xiao & Zhang, 2014) (and SVRG^{++} (Allen-Zhu & Yuan, 2016)), AccProxSVRG (Nitanda, 2014), Universal Catlyst (Lin et al., 2015) , APCG (Lin et al., 2014a), and Katyusha (Allen-Zhu, 2016). The details of the implemented algorithms and their parameter tunings are found in the supplementary material. In the experiments, we focus on the regularized logistic regression problem for binary classification, with regularizer $\lambda_1\|\cdot\|_1 + (\lambda_2/2)\|\cdot\|_2^2$.

We used three publicly available data sets in the experiments. Their sizes n and dimensions d, and common mini-batch sizes b for all implemented algorithms are listed in Table 2.

Data sets	n	d	b
a9a	$32,561$	123	180
rcv1	$20,242$	$47,236$	140
sido0	$12,678$	$4,932$	100

Table 2. Summary of the data sets and mini-batch size used in our numerical experiments

For regularization parameters, we used three settings $(\lambda_1, \lambda_2) = (10^{-4}, 0)$, $(10^{-4}, 10^{-6})$, and $(0, 10^{-6})$. For the former case, the objective is non-strongly convex, and for the latter two cases, the objectives are strongly convex. Note that for the latter two cases, the strong convexity of the objectives is $\mu = 10^{-6}$ and is relatively small; thus, it makes acceleration methods beneficial.

Figure 1 shows the comparisons of our method with the different methods described above on several settings. "Objective Gap" means $P(x) - P(x_*)$ for the output solution x. "Gradient Evaluations $/n$" is the number of computations of stochastic gradients ∇f_i divided by n. "Restart_DASVRDA" means DASVRDA with heuristic adaptive restarting. We can observe the following from these results:

- Our proposed DASVRDA and Restart_DASVRDA significantly outperformed all the other methods overall.
- DASVRDA with the heuristic adaptive restart scheme efficiently made use of the local strong convexities of non-strongly convex objectives and significantly outperformed vanilla DASVRDA on a9a and rcv1 data

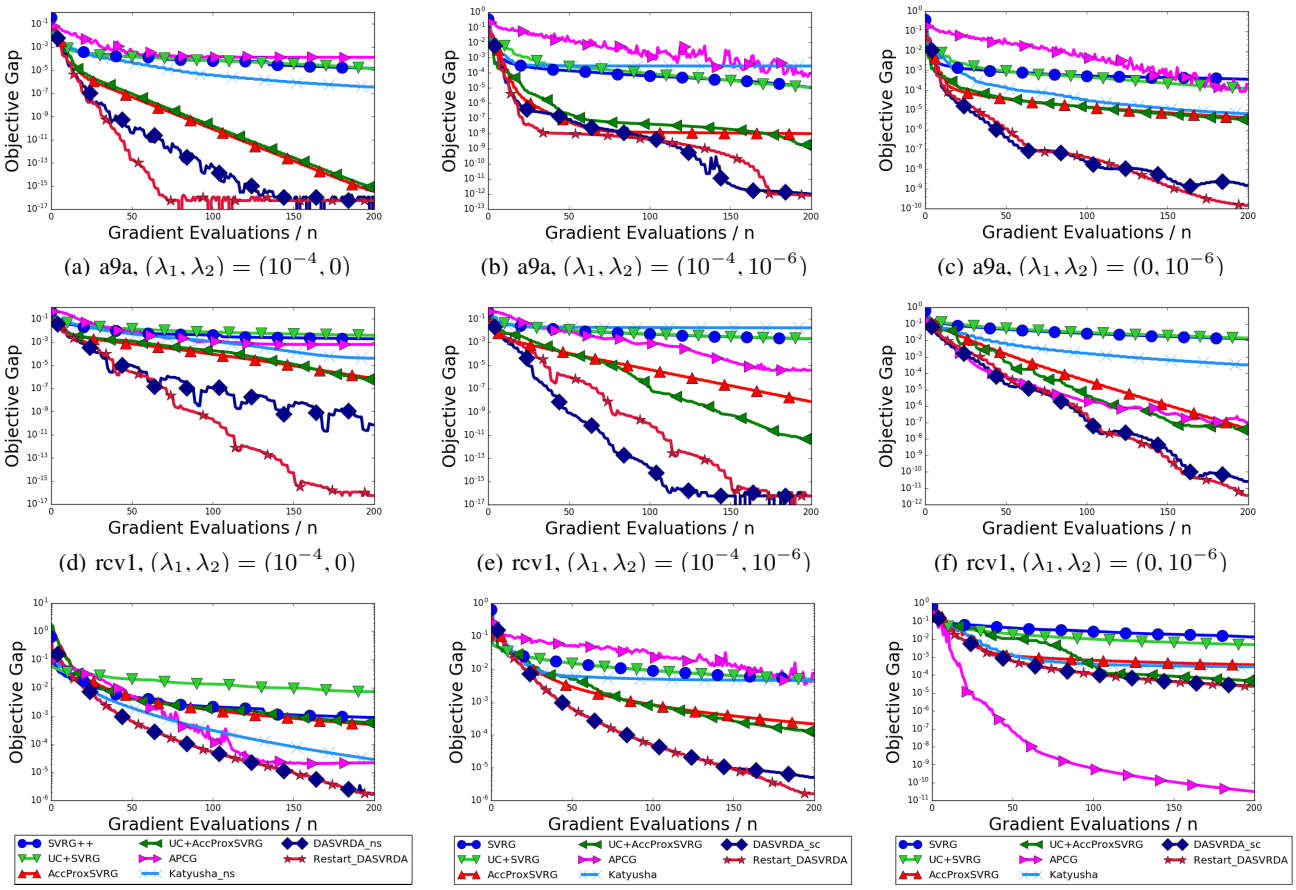

(a) a9a, $(\lambda_1, \lambda_2) = (10^{-4}, 0)$ (b) a9a, $(\lambda_1, \lambda_2) = (10^{-4}, 10^{-6})$ (c) a9a, $(\lambda_1, \lambda_2) = (0, 10^{-6})$

(d) rcv1, $(\lambda_1, \lambda_2) = (10^{-4}, 0)$ (e) rcv1, $(\lambda_1, \lambda_2) = (10^{-4}, 10^{-6})$ (f) rcv1, $(\lambda_1, \lambda_2) = (0, 10^{-6})$

(g) sido0, $(\lambda_1, \lambda_2) = (10^{-4}, 0)$ (h) sido0, $(\lambda_1, \lambda_2) = (10^{-4}, 10^{-6})$ (i) sido0, $(\lambda_1, \lambda_2) = (0, 10^{-6})$

Figure 1. Comparisons on a9a (top), rcv1 (middle) and sido0 (bottom) data sets, for regularization parameters $(\lambda_1, \lambda_2) = (10^{-4}, 0)$ (left), $(\lambda_1, \lambda_2) = (10^{-4}, 10^{-6})$ (middle) and $(\lambda_1, \lambda_2) = (0, 10^{-6})$ (right).

sets. For the other settings, the algorithm was still comparable to vanilla DASVRDA.

- UC+SVRG did not work as well as it did in theory, and its performances were almost the same as that of vanilla SVRG.
- UC+AccProxSVRG sometimes outperformed vanilla AccProxSVRG but was always outperformed by our methods.
- APCG sometimes performed unstably and was outperformed by vanilla SVRG. On sido0 data set, for Elastic Net Setting, APCG significantly outperformed all other methods.
- Katyusha outperformed vanilla SVRG overall. However, sometimes Katyusha was slower than vanilla SVRG for Elastic Net Settings. This is probably because SVRG is almost adaptive to local strong convexities of loss functions, whereas Katyusha is not (see the remark in supplementary material).

8. Conclusion

In this paper, we developed a new accelerated stochastic variance reduced gradient method for regularized empirical risk minimization problems in mini-batch settings: DASVRDA. We have shown that DASVRDA achieves the total computational costs of $O(d(n + \sqrt{nL/\varepsilon} + b\sqrt{L/\varepsilon}))$ and $O(d(n + \sqrt{nL/\mu} + b\sqrt{L/\mu})\log(1/\varepsilon))$ in size b mini-batch settings for non-strongly and strongly convex objectives, respectively. In addition, DASVRDA essentially achieves the optimal iteration complexities only with size $O(\sqrt{n})$ mini-batches for both settings. Furthermore, even in non-mini-batch settings, our method surpasses the best-known rate of $O(d(n\log(1/\varepsilon) + \sqrt{nL/\varepsilon}))$ for non-strongly convex objectives. In the numerical experiments, our method significantly outperformed state-of-the-art methods, including Katyusha and AccProxSVRG.

Acknowledgment

This work was partially supported by MEXT kakenhi (25730013, 25120012, 26280009, 15H01678 and 15H05707), JST-PRESTO and JST-CREST.

References

Agarwal, Alekh, Negahban, Sahand, and Wainwright, Martin J. Fast global convergence rates of gradient methods for high-dimensional statistical recovery. *Advances in Neural Information Processing Systems*, pp. 37–45, 2010.

Agarwal, Alekh, Negahban, Sahand, and Wainwright, Martin J. Stochastic optimization and sparse statistical recovery: Optimal algorithms for high dimensions. *Advances in Neural Information Processing Systems*, pp. 1538–1546, 2012.

Allen-Zhu, Zeyuan. Katyusha: The first direct acceleration of stochastic gradient methods. *ArXiv e-prints, abs/1603.05953*, 2016.

Allen-Zhu, Zeyuan and Yuan, Yang. Univr: A universal variance reduction framework for proximal stochastic gradient method. *arXiv preprint arXiv:1506.01972*, 2015.

Allen-Zhu, Zeyuan and Yuan, Yang. Improved SVRG for Non-Strongly-Convex or Sum-of-Non-Convex Objectives. In *Proceedings of the 33rd International Conference on Machine Learning*, ICML '16, 2016. Full version available at http://arxiv.org/abs/1506.01972.

Beck, Amir and Teboulle, Marc. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009.

Defazio, Aaron, Bach, Francis, and Lacoste-Julien, Simon. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. *Advances in Neural Information Processing Systems*, pp. 1646–1654, 2014.

Duchi, John, Hazan, Elad, and Singer, Yoram. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159, 2011.

Hazan, Elad, Agarwal, Amit, and Kale, Satyen. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.

Johnson, Rie and Zhang, Tong. Accelerating stochastic gradient descent using predictive variance reduction. *Advances in Neural Information Processing Systems*, pp. 315–323, 2013.

Konečnỳ, Jakub, Liu, Jie, Richtárik, Peter, and Takáč, Martin. Mini-batch semi-stochastic gradient descent in the proximal setting. *IEEE Journal of Selected Topics in Signal Processing*, 10(2):242–255, 2016.

Li, Huan and Lin, Zhouchen. Accelerated proximal gradient methods for nonconvex programming. *Advances in Neural Information Processing Systems*, pp. 379–387, 2015.

Lin, Hongzhou, Mairal, Julien, and Harchaoui, Zaid. A universal catalyst for first-order optimization. *Advances in Neural Information Processing Systems*, pp. 3384–3392, 2015.

Lin, Qihang, Lu, Zhaosong, and Xiao, Lin. An accelerated proximal coordinate gradient method. *Advances in Neural Information Processing Systems*, pp. 3059–3067, 2014a.

Lin, Qihang, Xiao, Lin, et al. An adaptive accelerated proximal gradient method and its homotopy continuation for sparse optimization. *ICML*, pp. 73–81, 2014b.

Nesterov, Yurii. Introductory lectures on convex optimization: A basic course. *Springer Science & Business Media*, 87, 2013.

Nesterov, Yurii et al. Gradient methods for minimizing composite objective function. Technical report, UCL, 2007.

Nitanda, Atsushi. Stochastic proximal gradient descent with acceleration techniques. *Advances in Neural Information Processing Systems*, pp. 1574–1582, 2014.

Nitanda, Atsushi. Accelerated stochastic gradient descent for minimizing finite sums. *arXiv preprint arXiv:1506.03016*, 2015.

O'Donoghue, Brendan and Candes, Emmanuel. Adaptive restart for accelerated gradient schemes. *Foundations of computational mathematics*, 15(3):715–732, 2015.

Roux, Nicolas L, Schmidt, Mark, and Bach, Francis R. A stochastic gradient method with an exponential convergence _rate for finite training sets. *Advances in Neural Information Processing Systems*, pp. 2663–2671, 2012.

Schmidt, Mark, Roux, Nicolas Le, and Bach, Francis. Minimizing finite sums with the stochastic average gradient. *arXiv preprint arXiv:1309.2388*, 2013.

Shalev-Shwartz, Shai and Singer, Yoram. Logarithmic regret algorithms for strongly convex repeated games. *The Hebrew University*, 2007.

Shalev-Shwartz, Shai and Zhang, Tong. Stochastic dual coordinate ascent methods for regularized loss. *The Journal of Machine Learning Research*, 14(1):567–599, 2013.

Singer, Yoram and Duchi, John C. Efficient learning using forward-backward splitting. *Advances in Neural Information Processing Systems*, pp. 495–503, 2009.

Tseng, Paul. On accelerated proximal gradient methods for convex-concave optimization. submitted to siam j. *J. Optim*, 2008.

Woodworth, Blake and Srebro, Nathan. Tight complexity bounds for optimizing composite objectives. *arXiv preprint arXiv:1605.08003*, 2016.

Xiao, Lin. Dual averaging method for regularized stochastic learning and online optimization. *Advances in Neural Information Processing Systems*, pp. 2116–2124, 2009.

Xiao, Lin and Zhang, Tong. A proximal stochastic gradient method with progressive variance reduction. *SIAM Journal on Optimization*, 24(4):2057–2075, 2014.

Zhang, Yuchen and Xiao, Lin. Stochastic primal-dual coordinate method for regularized empirical risk minimization. *Proceedings of the 32nd International Conference on Machine Learning*, 951:2015, 2015.

Supplementary material: Doubly Accelerated Stochastic Variance Reduced Dual Averaging Method for Regularized Empirical Risk Minimization

In this supplementary material, we give the proof of Theorem 5.1 and the optimality of γ_* (Section A), the lazy update algorithm of our method (Section B) and the experimental details (Section C).

A. Proof of Theorem 5.1

In this section, we give the comprehensive proof of Theorem 5.1. First we analyze One Stage Accelerated SVRDA algorithm.

Lemma A.1. *The sequence $\{\theta_k\}_{k\geq 1}$ defined in Algorithm 7 satisfies*

$$\theta_k - 1 = \theta_{k-2}$$

for $k \geq 1$, where $\theta_{-1} \overset{\text{def}}{=} 0$.

Proof. Since $\theta_k = \frac{k+1}{2}$ for $k \geq 0$, we have that

$$\theta_k - 1 = \frac{k+1}{2} - 1 = \frac{k-1}{2} = \theta_{k-2}.$$

□

Lemma A.2. *The sequence $\{\theta_k\}_{k\geq 1}$ defined in Algorithm 7 satisfies*

$$\theta_m \theta_{m-1} = \sum_{k=1}^{m} \theta_{k-1}.$$

□

Proof. Observe that

$$\theta_m \theta_{m-1} = \frac{m(m+1)}{4} = \sum_{k=1}^{m} \frac{k}{2} = \sum_{k=1}^{m} \theta_{k-1}.$$

□

Lemma A.3. *For every x, $y \in \mathbb{R}^d$,*

$$F(y) + \langle \nabla F(y), x - y \rangle + R(x) \leq P(x) - \frac{1}{2\bar{L}} \frac{1}{n} \sum_{i=1}^{n} \frac{1}{nq_i} \|\nabla f_i(x) - \nabla f_i(y)\|^2.$$

Proof. Since f_i is convex and L_i-smooth, we have (see (Nesterov, 2013))

$$f_i(y) + \langle \nabla f_i(y), x - y \rangle \leq f_i(x) - \frac{1}{2L_i} \|\nabla f_i(x) - \nabla f_i(y)\|^2.$$

By the definition of $\{q_i\}$, summing this inequality from $i = 1$ to n and dividing it by n results in

$$F(y) + \langle \nabla F(y), x - y \rangle \leq F(x) - \frac{1}{2\bar{L}} \frac{1}{n} \sum_{i=1}^{n} \frac{1}{nq_i} \|\nabla f_i(x) - \nabla f_i(y)\|^2.$$

Adding $R(x)$ to the both sides of this inequality gives the desired result. □

Lemma A.4.

$$\bar{g}_k = \frac{1}{\theta_k \theta_{k-1}} \sum_{k'=1}^{k} \theta_{k'-1} g_{k'} \ (k \geq 1).$$

Proof. For $k = 1$, $\bar{g}_1 = g_1 = \frac{1}{1 \cdot \frac{1}{2}} \sum_{k'=1}^{1} \frac{1}{2} \cdot g_{k'}$ by the definition of θ_0.
Assume that the claim holds for some $k \geq 1$. Then

$$\begin{aligned}
\bar{g}_{k+1} &= \left(1 - \frac{1}{\theta_{k+1}}\right) \bar{g}_k + \frac{1}{\theta_{k+1}} g_{k+1} \\
&= \left(1 - \frac{2}{k+2}\right) \frac{4}{(k+1)k} \sum_{k'=1}^{k} \theta_{k'-1} g_{k'} + \frac{2}{k+2} g_{k+1} \\
&= \frac{4}{(k+2)(k+1)} \sum_{k'=1}^{k+1} \theta_{k'-1} g_{k'} \\
&= \frac{1}{\theta_{k+1}\theta_k} \sum_{k'=1}^{k+1} \theta_{k'-1} g_{k'}.
\end{aligned}$$

The first equality follows from the definition of $\bar{g}_{k+1}$. Second equality is due to the assumption of the induction. This finishes the proof for Lemma A.4. □

Next we prove the following main lemma for One Stage Accelerated SVRDA. The proof is inspired by the one of AccSDA given in (Xiao, 2009).

Lemma A.5. *For One Stage Accelerated SVRDA, we have that*

$$\begin{aligned}
&\mathbb{E}[P(x_m) - P(x)] \\
&\leq \frac{2}{\eta(m+1)m} \|z_0 - x\|^2 - \frac{2}{\eta(m+1)m} \mathbb{E}\|z_m - x\|^2 \\
&\quad + \frac{2}{(m+1)m} \sum_{k=1}^{m} \left(\frac{(k+1)k\mathbb{E}\|g_k - \nabla F(y_k)\|^2}{4\left(\frac{1}{\eta} - \bar{L}\right)} - \frac{k}{2\bar{L}} \frac{1}{n} \sum_{i=1}^{n} \frac{1}{nq_i} \mathbb{E}\|\nabla f_i(y_k) - \nabla f_i(x)\|^2 \right),
\end{aligned}$$

for any $x \in \mathbb{R}^d$, where the expectations are taken with respect to $I_k (1 \leq k \leq m)$.

Proof. We define

$$\begin{aligned}
\ell_k(x) &= F(y_k) + \langle \nabla F(y_k), x - y_k \rangle + R(x), \\
\hat{\ell}_k(x) &= F(y_k) + \langle g_k, x - y_k \rangle + R(x).
\end{aligned}$$

Observe that $\ell_k, \hat{\ell}_k$ is convex and $\ell_k \leq P$ by the convexity of F and R. Moreover, for $k \geq 1$ we have that

$$\begin{aligned}
\sum_{k'=1}^{k} \theta_{k'-1} \hat{\ell}_{k'}(z) &= \sum_{k'=1}^{k} \theta_{k'-1} F(y_k) + \sum_{k'=1}^{k} \langle \theta_{k'-1} g_{k'}, z - y_{k'} \rangle + \sum_{k'=1}^{k} \theta_{k'-1} R(z) \\
&= \langle \theta_k \theta_{k-1} \bar{g}_k, z \rangle + \theta_k \theta_{k-1} R(z) + \sum_{k'=1}^{k} \theta_{k'-1} F(y_k) - \sum_{k'=1}^{k} \theta_{k'-1} \langle g_{k'}, y_{k'} \rangle.
\end{aligned}$$

The second equality follows from Lemma A.4 and Lemma A.2. Thus we see that $z_k = \operatorname{argmin}_{z \in \mathbb{R}^d} \left\{ \sum_{k'=1}^{k} \theta_{k'-1} \hat{\ell}_{k'}(z) + \frac{1}{2\eta} \|z - z_0\|^2 \right\}$. Observe that F is convex and $\bar{L}$-smooth. Thus we have that

$$F(x_k) \leq F(y_k) + \langle \nabla F(y_k), x_k - y_k \rangle + \frac{\bar{L}}{2} \|x_k - y_k\|^2. \tag{4}$$

Hence we see that

$$\begin{aligned}
P(x_k) &\leq \ell_k(x_k) + \frac{\bar{L}}{2}\|x_k - y_k\|^2 \\
&= \ell_k\left(\left(1-\frac{1}{\theta_k}\right)x_{k-1} + \frac{1}{\theta_k}z_k\right) + \frac{\bar{L}}{2}\left\|\left(1-\frac{1}{\theta_k}\right)x_{k-1} + \frac{1}{\theta_k}z_k - y_k\right\|^2 \\
&\leq \left(1-\frac{1}{\theta_k}\right)\ell_k(x_{k-1}) + \frac{1}{\theta_k}\ell_k(z_k) + \frac{\bar{L}}{2\theta_k^2}\|z_k - z_{k-1}\|^2 \\
&\leq \left(1-\frac{1}{\theta_k}\right)P(x_{k-1}) + \frac{1}{\theta_k\theta_{k-1}}\left(\theta_{k-1}\hat{\ell}_k(z_k) + \frac{\bar{L}}{2}\|z_k - z_{k-1}\|^2\right) \\
&\quad - \frac{1}{\theta_k}\langle g_k - \nabla F(y_k), z_k - y_k\rangle \\
&= \left(1-\frac{1}{\theta_k}\right)P(x_{k-1}) + \frac{1}{\theta_k\theta_{k-1}}\left(\theta_{k-1}\hat{\ell}_k(z_k) + \frac{1}{2\eta}\|z_k - z_{k-1}\|^2\right) \\
&\quad - \frac{1}{2\theta_k\theta_{k-1}}\left(\frac{1}{\eta} - \bar{L}\right)\|z_k - z_{k-1}\|^2 - \frac{1}{\theta_k}\langle g_k - \nabla F(y_k), z_k - z_{k-1}\rangle \\
&\quad - \frac{1}{\theta_k}\langle g_k - \nabla F(y_k), z_{k-1} - y_k\rangle.
\end{aligned}$$

The first inequality follows from (4). The first equality is due to the definition of x_k. The second inequality is due to the convexity of ℓ_k and the definition of y_k. The third inequality holds because $\ell_k \leq P$ and $\frac{1}{\theta_k^2} \leq \frac{1}{\theta_k\theta_{k-1}}$.

Since $\frac{1}{\eta} > \bar{L}$, we have that

$$\begin{aligned}
&-\frac{1}{2\theta_k\theta_{k-1}}\left(\frac{1}{\eta} - \bar{L}\right)\|z_k - z_{k-1}\|^2 - \frac{1}{\theta_k}\langle g_k - \nabla F(y_k), z_k - z_{k-1}\rangle \\
\leq &\frac{1}{\theta_k}\frac{\theta_{k-1}\|g_k - \nabla F(y_k)\|^2}{2\left(\frac{1}{\eta} - \bar{L}\right)} \\
\leq &\frac{\|g_k - \nabla F(y_k)\|^2}{2\left(\frac{1}{\eta} - \bar{L}\right)}.
\end{aligned}$$

The first inequality is due to Young's inequality. The second inequality holds because $\theta_{k-1} \leq \theta_k$.

Using this inequality, we get that

$$\begin{aligned}
P(x_k) \leq &\left(1-\frac{1}{\theta_k}\right)P(x_{k-1}) + \frac{1}{\theta_k\theta_{k-1}}\left(\theta_{k-1}\hat{\ell}_k(z_k) + \frac{1}{2\eta}\|z_k - z_{k-1}\|^2\right) \\
&+ \frac{\|g_k - \nabla F(y_k)\|^2}{2\left(\frac{1}{\eta} - \bar{L}\right)} - \frac{1}{\theta_k}\langle g_k - \nabla F(y_k), z_{k-1} - y_k\rangle.
\end{aligned}$$

Multiplying the both sides of the above inequality by $\theta_k\theta_{k-1}$ yields

$$\begin{aligned}
\theta_k\theta_{k-1}P(x_k) \leq &\theta_{k-1}(\theta_k - 1)P(x_{k-1}) + \theta_{k-1}\hat{\ell}_k(z_k) + \frac{1}{2\eta}\|z_k - z_{k-1}\|^2 \\
&+ \frac{\theta_k\theta_{k-1}\|g_k - \nabla F(y_k)\|^2}{2\left(\frac{1}{\eta} - \bar{L}\right)} - \theta_{k-1}\langle g_k - \nabla F(y_k), z_{k-1} - y_k\rangle.
\end{aligned}$$

By the fact that $\sum_{k'=1}^{k-1}\theta_{k'-1}\hat{\ell}_{k'}(z) + \frac{1}{2\eta}\|z - z_0\|^2$ is $\frac{1}{\eta}$-strongly convex and z_{k-1} is the minimizer of $\sum_{k'=1}^{k-1}\theta_{k'-1}\hat{\ell}_{k'}(z) + \frac{1}{2\eta}\|z - z_0\|^2$ for $k \geq 2$, we have that

$$\sum_{k'=1}^{k-1}\theta_{k'-1}\hat{\ell}_{k'}(z_{k-1}) + \frac{1}{2\eta}\|z_{k-1} - z_0\|^2 + \frac{1}{2\eta}\|z_k - z_{k-1}\|^2 \leq \sum_{k'=1}^{k-1}\theta_{k'-1}\hat{\ell}_{k'}(z_k) + \frac{1}{2\eta}\|z_k - z_0\|^2$$

for $k \geq 1$ (and, for $k = 1$, we define $\sum_{k'=1}^{0} = 0$).

Using this inequality, we obtain

$$\begin{aligned}
&\theta_k \theta_{k-1} P(x_k) - \sum_{k'=1}^{k} \theta_{k'-1} \hat{\ell}_{k'}(z_k) - \frac{1}{2\eta}\|z_k - z_0\|^2 \\
\leq &\theta_{k-1}\theta_{k-2} P(x_{k-1}) - \sum_{k'=1}^{k-1} \theta_{k'-1} \hat{\ell}_{k'}(z_{k-1}) - \frac{1}{2\eta}\|z_{k-1} - z_0\|^2 + \frac{\theta_k \theta_{k-1}}{2\left(\frac{1}{\eta} - \bar{L}\right)}\|g_k - \nabla F(y_k)\|^2 \\
&- \theta_{k-1}\langle g_k - \nabla F(y_k), z_{k-1} - y_k\rangle.
\end{aligned}$$

Here, the inequality follows from Lemma A.1 (we defined $\theta_{-1} \stackrel{\text{def}}{=} 0$).

Summing the above inequality from $k = 1$ to m results in

$$\begin{aligned}
&\theta_m \theta_{m-1} P(x_m) - \sum_{k=1}^{m} \theta_{k-1} \hat{\ell}_k(z_m) - \frac{1}{2\eta}\|z_m - z_0\|^2 \\
\leq &\sum_{k=1}^{m} \frac{\theta_k \theta_{k-1}\|g_k - \nabla F(y_k)\|^2}{2\left(\frac{1}{\eta} - \bar{L}\right)} - \sum_{k=1}^{m} \theta_{k-1}\langle g_k - \nabla F(y_k), z_{k-1} - y_k\rangle.
\end{aligned}$$

Using $\frac{1}{\eta}$-strongly convexity of the function $\sum_{k=1}^{m} \theta_{k-1}\hat{\ell}_k(z) + \frac{1}{2\eta}\|z - z_0\|^2$ and the optimality of z_m, we have that

$$\sum_{k=1}^{m} \theta_{k-1}\hat{\ell}_k(z_m) + \frac{1}{2\eta}\|z_m - z_0\|^2 \leq \sum_{k=1}^{m} \theta_{k-1}\hat{\ell}_k(x) + \frac{1}{2\eta}\|z_0 - x\|^2 - \frac{1}{2\eta}\|z_m - x\|^2.$$

From this inequality, we see that

$$\begin{aligned}
&\theta_m \theta_{m-1} P(x_m) \\
\leq &\sum_{k=1}^{m} \theta_{k-1}\hat{\ell}_k(x) + \frac{1}{2\eta}\|z_0 - x\|^2 - \frac{1}{2\eta}\|z_m - x\|^2 \\
&+ \sum_{k=1}^{m} \frac{\theta_k \theta_{k-1}\|g_k - \nabla F(y_k)\|^2}{2\left(\frac{1}{\eta} - \bar{L}\right)} - \sum_{k=1}^{m} \theta_{k-1}\langle g_k - \nabla F(y_k), z_{k-1} - y_k\rangle \\
= &\sum_{k=1}^{m} \theta_{k-1}\ell_k(x) + \frac{1}{2\eta}\|z_0 - x\|^2 - \frac{1}{2\eta}\|z_m - x\|^2 \\
&+ \sum_{k=1}^{m} \frac{\theta_k \theta_{k-1}\|g_k - \nabla F(y_k)\|^2}{2\left(\frac{1}{\eta} - \bar{L}\right)} - \sum_{k=1}^{m} \theta_{k-1}\langle g_k - \nabla F(y_k), z_{k-1} - x\rangle.
\end{aligned}$$

By Lemma A.3 with $x = x$ and $y = y_k$, we have that

$$\ell_k(x) \leq P(x) - \frac{1}{2\bar{L}}\frac{1}{n}\sum_{i=1}^{n} \frac{1}{nq_i}\|\nabla f_i(x) - \nabla f_i(y_k)\|^2.$$

Applying this inequality to the above inequality yields

$$\begin{aligned}
&\theta_m\theta_{m-1}P(x_m) - \sum_{k=1}^{m}\theta_{k-1}P(x) \\
\leq& \frac{1}{2\eta}\|z_0 - x\|^2 - \frac{1}{2\eta}\|z_m - x\|^2 \\
&+ \sum_{k=1}^{m}\left[\frac{\theta_k\theta_{k-1}\|g_k - \nabla F(y_k)\|^2}{2\left(\frac{1}{\eta} - \bar{L}\right)} - \frac{\theta_{k-1}}{2\bar{L}}\frac{1}{n}\sum_{i=1}^{n}\frac{1}{nq_i}\|\nabla f_i(x) - \nabla f_i(y_k)\|^2\right] \\
&- \sum_{k=1}^{m}\theta_{k-1}\langle g_k - \nabla F(y_k), z_{k-1} - x\rangle.
\end{aligned}$$

Using Lemma A.2 and dividing the both sides of the above inequality by $\theta_m\theta_{m-1}$ result in

$$\begin{aligned}
&P(x_m) - P(x) \\
\leq& \frac{1}{2\eta\theta_m\theta_{m-1}}\|z_0 - x\|^2 - \frac{1}{2\eta\theta_m\theta_{m-1}}\|z_m - x\|^2 \\
&+ \frac{1}{\theta_m\theta_{m-1}}\sum_{k=1}^{m}\left[\frac{\theta_k\theta_{k-1}\|g_k - \nabla F(y_k)\|^2}{2\left(\frac{1}{\eta} - \bar{L}\right)} - \frac{\theta_{k-1}}{2\bar{L}}\frac{1}{n}\sum_{i=1}^{n}\frac{1}{nq_i}\|\nabla f_i(x) - \nabla f_i(y_k)\|^2\right] \\
&- \frac{1}{\theta_m\theta_{m-1}}\sum_{k=1}^{m}\theta_{k-1}\langle g_k - \nabla F(y_k), z_{k-1} - x\rangle.
\end{aligned}$$

Taking the expectations with respect to $I_k (1 \leq k \leq m)$ on the both sides of this inequality yields

$$\begin{aligned}
&\mathbb{E}[P(x_m) - P(x)] \\
\leq& \frac{1}{2\eta\theta_m\theta_{m-1}}\|z_0 - x\|^2 - \frac{1}{2\eta\theta_m\theta_{m-1}}\mathbb{E}\|z_m - x\|^2 \\
&+ \frac{1}{\theta_m\theta_{m-1}}\sum_{k=1}^{m}\left[\frac{\theta_k\theta_{k-1}\mathbb{E}\|g_k - \nabla F(y_k)\|^2}{2\left(\frac{1}{\eta} - \bar{L}\right)} - \frac{\theta_{k-1}}{2\bar{L}}\frac{1}{n}\sum_{i=1}^{n}\frac{1}{nq_i}\mathbb{E}\|\nabla f_i(x) - \nabla f_i(y_k)\|^2\right].
\end{aligned}$$

Here we used the fact that $\mathbb{E}[g_k - \nabla F(y_k)] = 0$ for $k = 1, \ldots, m$. This finishes the proof of Lemma A.5. □

Now we need the following lemma.

Lemma A.6. *For every $x \in \mathbb{R}^d$,*

$$\frac{1}{n}\sum_{i=1}^{n}\frac{1}{nq_i}\|\nabla f_i(x) - \nabla f_i(x_*)\|^2 \leq 2\bar{L}(P(x) - P(x_*)).$$

Proof. From the argument of the proof of Lemma A.3, we have

$$\frac{1}{n}\sum_{i=1}^{n}\frac{1}{nq_i}\|\nabla f_i(x) - \nabla f_i(x_*)\|^2 \leq 2\bar{L}(F(x) - \langle\nabla F(x_*), x - x_*\rangle - F(x_*)).$$

By the optimality of x_*, there exists $\xi_* \in \partial R(x_*)$ such that $\nabla F(x_*) + \xi_* = 0$. Then we have

$$-\langle\nabla F(x_*), x - x_*\rangle = \langle\xi_*, x - x_*\rangle \leq R(x) - R(x_*),$$

and hence

$$\frac{1}{n}\sum_{i=1}^{n}\frac{1}{nq_i}\|\nabla f_i(x) - \nabla f_i(x_*)\|^2 \leq 2\bar{L}(P(x) - P(x_*)).$$

□

Proposition A.7. *Let $\gamma > 1$. For One Pass Accelerated SVRDA, it follows that*

$$\mathbb{E}[P(x_m) - P(\widetilde{x})] \leq \frac{2}{\eta(m+1)m}\|\widetilde{y} - \widetilde{x}\|^2 - \frac{2}{\eta(m+1)m}\mathbb{E}\|z_m - \widetilde{x}\|^2,$$

and

$$\begin{aligned}&\mathbb{E}[P(x_m) - P(x_*)]\\ &\leq \frac{1}{\gamma}(P(\widetilde{x}) - P(x_*)) + \frac{2}{\eta(m+1)m}\|\widetilde{y} - x_*\|^2 - \frac{2}{\eta(m+1)m}\mathbb{E}\|z_m - x_*\|^2,\end{aligned}$$

where the expectations are taken with respect to $I_k (1 \leq k \leq m)$.

Proof. We bound the variance of the averaged stochastic gradient $\mathbb{E}\|g_k - \nabla F(y_k)\|^2$:

$$\begin{aligned}&\mathbb{E}\|g_k - \nabla F(y_k)\|^2\\ &= \mathbb{E}\left[\mathbb{E}_{I_k}\|g_k - \nabla F(y_k)\|^2 \mid [k-1]\right]\\ &= \frac{1}{b}\mathbb{E}\left[\mathbb{E}_{i\sim Q}\|(\nabla f_i(y_k) - \nabla f_i(\widetilde{x}))/nq_i + \nabla F(\widetilde{x}) - \nabla F(y_k)\|^2 \mid [k-1]\right]\\ &\leq \frac{1}{b}\mathbb{E}\left[\mathbb{E}_{i\sim Q}\|(\nabla f_i(y_k) - \nabla f_i(\widetilde{x}))/nq_i\|^2 \mid [k-1]\right]\\ &= \frac{1}{b}\mathbb{E}\left[\frac{1}{n}\sum_{i=1}^n \frac{1}{nq_i}\|\nabla f_i(y_k) - \nabla f_i(\widetilde{x})\|^2\right] \qquad (5)\\ &\leq \frac{2}{b}\mathbb{E}\left[\frac{1}{n}\sum_{i=1}^n \frac{1}{nq_i}\|\nabla f_i(y_k) - \nabla f_i(x_*)\|^2\right]\\ &\quad + \frac{2}{b}\mathbb{E}\left[\frac{1}{n}\sum_{i=1}^n \frac{1}{nq_i}\|\nabla f_i(\widetilde{x}) - \nabla f_i(x_*)\|^2\right]\\ &\leq \frac{2}{b}\mathbb{E}\left[\frac{1}{n}\sum_{i=1}^n \frac{1}{nq_i}\|\nabla f_i(y_k) - \nabla f_i(x_*)\|^2\right] + \frac{4\bar{L}}{b}(P(\widetilde{x}) - P(x_*)). \qquad (6)\end{aligned}$$

The second equality follows from the independency of the random variables $\{i_1, \ldots, i_b\}$ and the unbiasedness of $(\nabla f_i(y_k) - \nabla f_i(\widetilde{x}))/nq_i + \nabla F(\widetilde{x})$. The first inequality is due to the fact that $\mathbb{E}\|X - \mathbb{E}[X]\|^2 \leq \mathbb{E}\|X\|^2$. The second inequality follows from Young's inequality. The final inequality is due to Lemma A.6.

Since $\frac{1}{\eta} = \left(1 + \frac{\gamma(m+1)}{b}\right)\bar{L}$ and $\gamma > 1$, using (5) yields

$$\frac{(k+1)k}{4\left(\frac{1}{\eta} - \bar{L}\right)}\mathbb{E}\|g_k - \nabla F(y_k)\|^2 - \frac{k}{2\bar{L}}\mathbb{E}\left[\frac{1}{n}\sum_{i=1}^n \frac{1}{nq_i}\|\nabla f_i(y_k) - \nabla f_i(\widetilde{x})\|^2\right] \leq 0.$$

By Lemma A.5 (with $x = \widetilde{x}$) we have

$$\mathbb{E}[P(x_m) - P(\widetilde{x})] \leq \frac{2}{\eta(m+1)m}\|\widetilde{y} - \widetilde{x}\|^2 - \frac{2}{\eta(m+1)m}\mathbb{E}\|z_m - \widetilde{x}\|^2.$$

Similarly, combining Lemma A.5 (with $x = x_*$) with (6) results in

$$\begin{aligned}&\mathbb{E}[P(x_m) - P(x_*)]\\ &\leq \frac{1}{\gamma}(P(\widetilde{x}) - P(x_*)) + \frac{2}{\eta(m+1)m}\|\widetilde{y} - x_*\|^2 - \frac{2}{\eta(m+1)m}\mathbb{E}\|z_m - x_*\|^2.\end{aligned}$$

These are the desired results. □

Lemma A.8. *The sequence $\{\widetilde{\theta}_s\}_{s\geq 1}$ defined in Algorithm 6 satisfies*

$$\widetilde{\theta}_s\left(\widetilde{\theta}_s - 1 + \frac{1}{\gamma}\right) \leq \widetilde{\theta}_{s-1}^2$$

for any $s \geq 1$.

Proof. Since $\widetilde{\theta}_s = \left(1 - \frac{1}{\gamma}\right)\frac{s+1}{2}$ for $s \geq 1$, we have

$$\begin{aligned}
&\widetilde{\theta}_s\left(\widetilde{\theta}_s - 1 + \frac{1}{\gamma}\right) \\
&= \left(1 - \frac{1}{\gamma}\right)\frac{s+1}{2}\left(\left(1 - \frac{1}{\gamma}\right)\frac{s+1}{2} - 1 + \frac{1}{\gamma}\right) \\
&= \left(1 - \frac{1}{\gamma}\right)^2 \frac{s^2-1}{4} \\
&\leq \widetilde{\theta}_{s-1}^2.
\end{aligned}$$

This finishes the proof of Lemma A.8. □

Now we are ready to proof Theorem 5.1.

Proof of Theorem 5.1. By Proposition A.7, we have

$$\mathbb{E}[P(\widetilde{x}_s) - P(\widetilde{x}_{s-1})] \leq \frac{2}{\eta(m+1)m}\mathbb{E}\|\widetilde{y}_s - \widetilde{x}_{s-1}\|^2 - \frac{2}{\eta(m+1)m}\mathbb{E}\|\widetilde{z}_s - \widetilde{x}_{s-1}\|^2,$$

and

$$\begin{aligned}
&\mathbb{E}[P(\widetilde{x}_s) - P(x_*)] \\
&\leq \frac{1}{\gamma}\mathbb{E}[P(\widetilde{x}_{s-1}) - P(x_*)] + \frac{2}{\eta(m+1)m}\mathbb{E}\|\widetilde{y}_s - x_*\|^2 - \frac{2}{\eta(m+1)m}\mathbb{E}\|\widetilde{z}_s - x_*\|^2,
\end{aligned}$$

where the expectations are taken with respect to the history of all random variables.

Hence we have

$$\mathbb{E}[P(\widetilde{x}_s) - P(\widetilde{x}_{s-1})] \leq \frac{4}{\eta(m+1)m}\mathbb{E}\langle \widetilde{z}_s - \widetilde{y}_s, \widetilde{x}_{s-1} - \widetilde{y}_s\rangle - \frac{2}{\eta(m+1)m}\mathbb{E}\|\widetilde{z}_s - \widetilde{y}_s\|^2, \tag{7}$$

and

$$\begin{aligned}
&\mathbb{E}[P(\widetilde{x}_s) - P(x_*)] \\
&\leq \frac{1}{\gamma}\mathbb{E}[P(\widetilde{x}_{s-1}) - P(x_*)] + \frac{4}{\eta(m+1)m}\mathbb{E}\langle \widetilde{z}_s - \widetilde{y}_s, x_* - \widetilde{y}_s\rangle - \frac{2}{\eta(m+1)m}\mathbb{E}\|\widetilde{z}_s - \widetilde{y}_s\|^2.
\end{aligned} \tag{8}$$

Multiplying (7) by $\widetilde{\theta}_s(\widetilde{\theta}_s - 1)$ and adding $\widetilde{\theta}_s\times$ (8) yield

$$\begin{aligned}
&\widetilde{\theta}_s^2\mathbb{E}[P(\widetilde{x}_s) - P(x_*)] - \widetilde{\theta}_s\left(\widetilde{\theta}_s - 1 + \frac{1}{\gamma}\right)\mathbb{E}[P(\widetilde{x}_{s-1}) - P(x_*)] \\
&\leq \frac{4}{\eta(m+1)m}\mathbb{E}\langle \widetilde{\theta}_s(\widetilde{z}_s - \widetilde{y}_s), (\widetilde{\theta}_s - 1)\widetilde{x}_{s-1} - \widetilde{\theta}_s\widetilde{y}_s + x_*\rangle - \frac{2}{\eta(m+1)m}\mathbb{E}\|\widetilde{\theta}_s(\widetilde{z}_s - \widetilde{y}_s)\|^2.
\end{aligned}$$

By Lemma (A.8), we have

$$\widetilde{\theta}_s\left(\widetilde{\theta}_s - 1 + \frac{1}{\gamma}\right) \leq \widetilde{\theta}_{s-1}^2$$

for $s \geq 1$.

Thus we get

$$\begin{aligned}
&\widetilde{\theta}_s^2\mathbb{E}[P(\widetilde{x}_s) - P(x_*)] - \widetilde{\theta}_{s-1}^2\mathbb{E}[P(\widetilde{x}_{s-1}) - P(x_*)] \\
&\leq \frac{4}{\eta(m+1)m}\mathbb{E}\langle\widetilde{\theta}_s(\widetilde{z}_s - \widetilde{y}_s), (\widetilde{\theta}_s - 1)\widetilde{x}_{s-1} - \widetilde{\theta}_s\widetilde{y}_s + x_*\rangle - \frac{2}{\eta(m+1)m}\mathbb{E}\|\widetilde{\theta}_s(\widetilde{z}_s - \widetilde{y}_s)\|^2. \\
&= \frac{2}{\eta(m+1)m}\left(\mathbb{E}\|(\widetilde{\theta}_s - 1)\widetilde{x}_{s-1} - \widetilde{\theta}_s\widetilde{y}_s + x_*\|^2 - \mathbb{E}\|(\widetilde{\theta}_s - 1)\widetilde{x}_{s-1} - \widetilde{\theta}_s\widetilde{z}_s + x_*\|^2\right)
\end{aligned}$$

Since $\widetilde{y}_s = \widetilde{x}_{s-1} + \frac{\widetilde{\theta}_{s-1}-1}{\widetilde{\theta}_s}(\widetilde{x}_{s-1} - \widetilde{x}_{s-2}) + \frac{\widetilde{\theta}_{s-1}}{\widetilde{\theta}_s}(\widetilde{z}_{s-1} - \widetilde{x}_{s-1})$, we have

$$(\widetilde{\theta}_s - 1)\widetilde{x}_{s-1} - \widetilde{\theta}_s\widetilde{y}_s + x_* = (\widetilde{\theta}_{s-1} - 1)\widetilde{x}_{s-2} - \widetilde{\theta}_{s-1}\widetilde{z}_{s-1} + x_*.$$

Therefore summing the above inequality from $s = 1$ to S, we obtain

$$\begin{aligned}
&\widetilde{\theta}_s^2\mathbb{E}[P(\widetilde{x}_S) - P(x_*)] \\
&\leq \frac{2}{\eta(m+1)m}\|(\widetilde{\theta}_0 - 1)\widetilde{x}_{-1} - \widetilde{\theta}_0\widetilde{z}_0 + x_*\|^2 \\
&= \frac{2}{\eta(m+1)m}\|\widetilde{x}_0 - x_*\|^2.
\end{aligned}$$

Dividing both sides by $\widetilde{\theta}_s^2$ finishes the proof of Theorem 5.1. □

Optimal choice of γ

We can choose the optimal value of γ based on the following lemma.

Lemma A.9. *Define* $g(\gamma) \overset{\text{def}}{=} \frac{\left(1+\frac{\gamma(m+1)}{b}\right)}{\left(1-\frac{1}{\gamma}\right)^2}$ *for* $\gamma > 1$. *Then,*

$$\gamma_* \overset{\text{def}}{=} \operatorname*{argmin}_{\gamma>1} g(\gamma) = \frac{1}{2}\left(3 + \sqrt{9 + \frac{8b}{m+1}}\right).$$

Proof. First observe that

$$g'(\gamma) = \frac{\frac{m+1}{b}\left(1 - \frac{1}{\gamma}\right)^2 - 2\left(1 + \frac{\gamma(m+1)}{b}\right)\left(1 - \frac{1}{\gamma}\right)\frac{1}{\gamma^2}}{\left(1 - \frac{1}{\gamma}\right)^2}.$$

Hence we have

$$\begin{aligned}
&g'(\gamma) = 0 \\
\iff &\frac{m+1}{b}\left(1 - \frac{1}{\gamma}\right)^2 - 2\left(1 + \frac{\gamma(m+1)}{b}\right)\left(1 - \frac{1}{\gamma}\right)\frac{1}{\gamma^2} = 0 \\
\iff &\frac{m+1}{b}(\gamma^2 - \gamma) - 2\left(1 + \frac{\gamma(m+1)}{b}\right) = 0 \\
\iff &\gamma^2 - 3\gamma - \frac{2b}{m+1} = 0 \\
\iff &\gamma = \frac{1}{2}\left(3 + \sqrt{9 + \frac{8b}{m+1}}\right) = \gamma_*.
\end{aligned}$$

Here the second and last equivalencies hold from $\gamma > 1$. Moreover observe that $g'(\gamma) > 0$ for $\gamma > \gamma_*$ and $g'(\gamma) < 0$ for $1 < \gamma < \gamma_*$. This means that γ_* is the minimizer of g on the region $\gamma > 1$. □

B. Lazy Update Algorithm of DASVRDA Method

In this section, we discuss how to efficiently compute the updates of the DASVRDA algorithm for sparse data. Specifically, we derive lazy update rules of One Stage Accelerated SVRDA for the following empirical risk minimization problem:

$$\frac{1}{n}\sum_{i=1}^{n}\psi_i(a_i^\top x)+\lambda_1\|x\|_1+\frac{\lambda_2}{2}\|x\|_2^2,\quad \lambda_1,\lambda_2\geq 0$$

For the sake of simplicity, we define the one dimensional soft-thresholding operator as follows:

$$\mathrm{soft}(z,\lambda)\stackrel{\mathrm{def}}{=}\mathrm{sign}\,(z)\max\{|z|-\lambda,0\},$$

for $z\in\mathbb{R}$. Moreover, in this section, we denote $[z_1,z_2]$ as $\{z\in\mathbb{Z}\mid z_1\leq z\leq z_2\}$ for integers $z_1,z_2\in\mathbb{Z}$. The explicit algorithm of the lazy updates for One Stage Accelerated SVRDA is given by Algorithm 9. Let us analyze the iteration cost of the algorithm. Suppose that each feature vector a_i is sparse and the expected number of the nonzero elements is $O(d')$. First note that $|\mathcal{A}_k|=O(bd')$ expectedly if $d'\ll d$. For updating x_{k-1}, by Proposition B.1, we need to compute $\sum_{k'\in K_j^\pm}\theta_{k'-2}/(1+\eta\theta_{k'-1}\theta_{k'-2}\lambda_2)$ and $\sum_{k'\in K_j^\pm}\theta_{k'-1}\theta_{k'-2}^2/(1+\eta\theta_{k'-1}\theta_{k'-2}\lambda_2)$ for each $j\in\mathcal{A}_k$. For this, we first make lists $\{S_k\}_{k=1}^m=\{\sum_{k'=1}^k\theta_{k'-2}/(1+\eta\theta_{k'-1}\theta_{k'-2}\lambda_2)\}_{k=1}^m$ and $\{S'_k\}_{k=1}^m=\{\sum_{k'=1}^k\theta_{k'-1}\theta_{k'-2}^2/(1+\eta\theta_{k'-1}\theta_{k'-2}\lambda_2)\}_{k=1}^m$ before running the algorithm. This needs only $O(m)$. Note that these lists are not depend on coordinate j. Since $K_j^\pm$ are sets of continuous integers in $[k_j+2,k]$ or unions of two sets of continuous integers in $[k_j+2,k]$, we can efficiently compute the above sums. For example, if $K_j^+=[k_j+2,s_-]\cup[s_+,k]$ for some integers $s_\pm\in[k_j+2,k]$, we can compute $\sum_{k'\in K_j^+}\theta_{k'-2}/(1+\eta\theta_{k'-1}\theta_{k'-2}\lambda_2)$ as $S_{s_-}-S_{k_j+1}+S_k-S_{s_+-1}$ and this costs only $O(1)$. Thus, for computing x_{k-1} and y_k, we need only $O(bd')$ computational cost. For computing g_k, we need to compute the inner product $a_i^\top y_k$ for each $i\in I_k$ and this costs $O(bd')$ expectedly. The expected cost of the rest of the updates is apparently $O(bd')$. Hence, the total expected iteration cost of our algorithm in serial settings becomes $O(bd')$ rather than $O(bd)$. Furthermore, we can extend our algorithm to parallel computing settings. Indeed, if we have b processors, processor b' runs on the set $\mathcal{A}_k^{b'}\stackrel{\mathrm{def}}{=}\{j\in[d]\mid a_{i_{b'},j}\neq 0\}$. Then the total iteration cost per processor becomes ideally $O(d')$. Generally the overlap among the sets $\mathcal{A}_k^{b'}$ may cause latency, however for sufficiently sparse data, this latency is negligible. The following proposition guarantees that Algorithm 9 is equivalent to Algorithm 7 when $R(x)=\lambda_1\|x\|_1+(\lambda_2/2)\|x\|_2^2$.

Proposition B.1. *Suppose that $R(x)=\lambda_1\|x\|_1+\frac{\lambda_2}{2}\|x\|_2^2$ with $\lambda_1,\lambda_2\geq 0$. Let $j\in[d]$, $k_j\in[m]\cup\{0\}$ and $k\geq k_j+1$. Assume that $\nabla_j f_i(y_{k'})=\nabla_j f_i(\widetilde{x})=0$ for any $i\in[b]$ and $k'\in[k_j+1,k-1]$. In Algorithm 7, the following results hold:*

$$x_{k-1,j}=\begin{cases}x_{0,j} & (k=1)\\ \frac{\theta_{k_j}\theta_{k_j-1}}{\theta_{k-1}\theta_{k-2}}x_{k_j,j}+\frac{1}{\theta_{k-1}\theta_{k-2}}\sum_{k'\in K_j^+}\frac{\theta_{k'-2}}{1+\eta\theta_{k'-1}\theta_{k'-2}\lambda_2}(z_{0,j}-M_{k',j}^+) & (k\geq 2)\\ \qquad+\frac{1}{\theta_{k-1}\theta_{k-2}}\sum_{k'\in K_j^-}\frac{\theta_{k'-2}}{1+\eta\theta_{k'-1}\theta_{k'-2}\lambda_2}(z_{0,j}-M_{k',j}^-)\end{cases},$$

$$y_{k,j}=\begin{cases}x_{0,j} & (k=1)\\ \left(1-\frac{1}{\theta_k}\right)x_{k-1,j}+\frac{1}{\theta_k}\frac{1}{1+\eta\theta_{k-1}\theta_{k-2}\lambda_2}\times & (k\geq 2)\\ \quad\mathrm{soft}\left(z_{0,j}-\eta g_{k_j,j}^{\mathrm{sum}}-\eta(\theta_{k-1}\theta_{k-2}-\theta_{k_j}\theta_{k_j-1})\widetilde{\nabla}_j,\eta\theta_{k-1}\theta_{k-2}\lambda_1\right)\end{cases},$$

and

$$z_{k,j}=\frac{1}{1+\eta\theta_k\theta_{k-1}\lambda_2}\mathrm{soft}(z_{0,j}-\eta g_{k,j}^{\mathrm{sum}},\eta\theta_k\theta_{k-1}\lambda_1),$$

where

$$M_{k',j}^\pm\stackrel{\mathrm{def}}{=}\eta\theta_{k'-1}\theta_{k'-2}(\widetilde{\nabla}_j\pm\lambda_1)+\eta g_{k_j,j}^{\mathrm{sum}}-\eta\theta_{k_j}\theta_{k_j-1}\widetilde{\nabla}_j,$$

and $K_j^\pm\subset[k_j+2,k]$ are defined as follows:

Let $c_1\stackrel{\mathrm{def}}{=}\frac{\eta\widetilde{\nabla}_j}{4}$, $c_2\stackrel{\mathrm{def}}{=}\frac{\eta\lambda_1}{4}$ and $c_3\stackrel{\mathrm{def}}{=}\eta g_{k_j,j}^{\mathrm{sum}}-\eta\theta_{k_j}\theta_{k_j-1}\widetilde{\nabla}_j$ to simplify the notation. Note that $c_2\geq 0$. Moreover, we

Algorithm 9 Lazy Updates for One Stage AccSVRDA $(\widetilde{y}, \widetilde{x}, \eta, m, b, Q)$

Input: $\widetilde{y}, \widetilde{x}, \eta > 0$, $m \in \mathbb{N}$, $b \in [n]$, Q.
  $x_0 = z_0 = \widetilde{y}$.
  $g_{0,j}^{\text{sum}} = 0$ $(j \in [d])$.
  $\theta_0 = \frac{1}{2}$.
  $k_j = 0$ $(j \in [d])$.
  $\widetilde{\nabla} = \nabla F(\widetilde{x})$.
  for $k = 1$ to m **do**
    Sample independently $i_1, \ldots, i_b \sim Q$. $I_k = \{i_1, \ldots, i_b\}$.
    $\mathcal{A}_k = \{j \in [d] \mid \exists b' \in [b] : a_{i_{b'},j} \neq 0\}$.
    $\theta_k = \frac{k+1}{2}$.
    for $j \in \mathcal{A}_k$ **do**
      Update $x_{k-1,j}$, $y_{k,j}$ as in Proposition B.1.
    end for
    for $j \in \mathcal{A}_k$ **do**
      $g_{k,j} = \frac{1}{b}\sum_{i \in I_t} \frac{1}{nq_i}\left(\psi_i'(a_i^\top y_k)a_{i,j} - \psi_i'(a_i^\top \widetilde{x})a_{i,j}\right) + \widetilde{\nabla}_j$.
      $g_{k,j}^{\text{sum}} = g_{k_j,j}^{\text{sum}} + \theta_{k-1} g_{k,j} + \left(\theta_k\theta_{k-1} - \theta_{k_j}\theta_{k_j-1}\right)\widetilde{\nabla}_j$.
      $z_{k,j} = \frac{1}{1+\eta\theta_k\theta_{k-1}\lambda_2}\text{soft}(z_{0,j} - \eta g_{k,j}^{\text{sum}}, \eta\theta_k\theta_{k-1}\lambda_1)$.
      $x_{k,j} = \left(1 - \frac{1}{\theta_k}\right)x_{k-1,j} + \frac{1}{\theta_k} z_{k,j}$.
      $k_j = k$.
    end for
  end for
Output: (x_m, z_m).

define

$$
\begin{aligned}
D_\pm &\overset{\text{def}}{=} (c_1 \pm c_2)^2 + 4(c_1 \pm c_2)(z_{0,j} - c_3), \\
s_\pm^+ &\overset{\text{def}}{=} \frac{c_1 + c_2 \pm \sqrt{D_+}}{c_1 + c_2}, \\
s_\pm^- &\overset{\text{def}}{=} \frac{c_1 - c_2 \pm \sqrt{D_-}}{c_1 - c_2},
\end{aligned}
$$

where if $s_\pm^\pm$ are not well defined, we simply assign 0 (or any number) to $s_\pm^\pm$.

1) *If $c_1 > c_2$, then*

$$
\begin{aligned}
K_j^+ &\overset{\text{def}}{=} \begin{cases} \emptyset & (D_+ \leq 0) \\ [k_j + 2, k] \cap [\lceil s_-^+ \rceil, \lfloor s_+^+ \rfloor] & (D_+ > 0) \end{cases}, \\
K_j^- &\overset{\text{def}}{=} \begin{cases} [k_j + 2, k] & (D_- \leq 0) \\ [k_j + 2, \lfloor s_-^- \rfloor] \cup [\lceil s_+^- \rceil, k] & (D_- > 0) \end{cases}.
\end{aligned}
$$

2) *If $c_1 = c_2$, then*

$$
\begin{aligned}
K_j^+ &\overset{\text{def}}{=} \begin{cases} \emptyset & (c_2 = 0 \wedge z_{0,j} \leq c_3) \\ [k_j + 2, k] & (c_2 = 0 \wedge z_{0,j} > c_3) \\ \emptyset & (c_2 > 0 \wedge D_+ \leq 0) \\ [k_j + 2, k] \cap [\lceil s_-^+ \rceil, \lfloor s_+^+ \rfloor] & (c_2 > 0 \wedge D_+ > 0) \end{cases}, \\
K_j^- &\overset{\text{def}}{=} \begin{cases} [k_j + 2, k] & (z_{0,j} < c_3) \\ \emptyset & (z_{0,j} \geq c_3) \end{cases}.
\end{aligned}
$$

3) *If* $|c_1| < c_2$, *then*

$$K_j^+ \stackrel{\text{def}}{=} \begin{cases} \emptyset & (D_+ \leq 0) \\ [k_j+2, k] \cap [\lceil s_-^+ \rceil, \lfloor s_+^+ \rfloor] & (D_+ > 0) \end{cases},$$
$$K_j^- \stackrel{\text{def}}{=} \begin{cases} \emptyset & (D_- \leq 0) \\ [k_j+2, k] \cap [\lceil s_-^- \rceil, \lfloor s_+^- \rfloor] & (D_- > 0) \end{cases}.$$

4) *If* $c_1 = -c_2$, *then*

$$K_j^+ \stackrel{\text{def}}{=} \begin{cases} \emptyset & (z_{0,j} \leq c_3) \\ [k_j+2, k] & (z_{0,j} > c_3) \end{cases}$$
$$K_j^- \stackrel{\text{def}}{=} \begin{cases} [k_j+2, k] & (c_2 = 0 \wedge z_{0,j} < c_3) \\ \emptyset & (c_2 = 0 \wedge z_{0,j} \geq c_3) \\ \emptyset & (c_2 > 0 \wedge D_- \leq 0) \\ [k_j+2, k] \cap [\lceil s_-^- \rceil, \lfloor s_+^- \rfloor] & (c_2 > 0 \wedge D_- > 0) \end{cases}.$$

5) *If* $c_1 < -c_2$, *then*

$$K_j^+ \stackrel{\text{def}}{=} \begin{cases} [k_j+2, k] & (D_+ \leq 0) \\ [k_j+2, \lfloor s_-^+ \rfloor] \cup [\lceil s_+^+ \rceil, k] & (D_+ > 0) \end{cases},$$
$$K_j^- \stackrel{\text{def}}{=} \begin{cases} \emptyset & (D_- \leq 0) \\ [k_j+2, k] \cap [\lceil s_-^- \rceil, \lfloor s_+^- \rfloor] & (D_- > 0) \end{cases}.$$

Proof. First we consider the case $k = 1$. Observe that

$$y_{1,j} = \left(1 - \frac{1}{\theta_1}\right) x_{0,j} + \frac{1}{\theta_1} z_{0,j} = z_{0,j} = x_{0,j},$$

and

$$\begin{aligned} z_{1,j} &= \frac{1}{1+\eta\theta_1\theta_0\lambda_2} \mathrm{soft}\left(z_{0,j} - \eta\theta_1\theta_0 \frac{1}{\theta_1} g_{1,j}, \eta\theta_1\theta_0\lambda_1\right) \\ &= \frac{1}{1+\eta\theta_1\theta_0\lambda_2} \mathrm{soft}\left(z_{0,j} - \eta\theta_0 g_{1,j}, \eta\theta_1\theta_0\lambda_1\right) \\ &= \frac{1}{1+\eta\theta_1\theta_0\lambda_2} \mathrm{soft}\left(z_{0,j} - \eta g_{1,j}^{\mathrm{sum}}, \eta\theta_1\theta_0\lambda_1\right). \end{aligned}$$

Next we consider the case $k \geq 2$. We show that

$$x_{k-1,j} = \frac{\theta_{k_j}\theta_{k_j-1}}{\theta_{k-1}\theta_{k-2}} x_{k_j,j} + \frac{1}{\theta_{k-1}\theta_{k-2}} \sum_{k'=k_j+2}^{k} \theta_{k'-2} z_{k'-1}. \tag{9}$$

For $k = k_j + 1$, (9) holds. Assume that (9) holds for some $k' \geq k_j + 1$. Then

$$\begin{aligned} x_{k',j} &= \left(1 - \frac{1}{\theta_{k'}}\right) x_{k'-1,j} + \frac{1}{\theta_{k'}} z_{k',j} \\ &= \left(1 - \frac{1}{\theta_{k'}}\right) \frac{\theta_{k_j}\theta_{k_j-1}}{\theta_{k'-1}\theta_{k'-2}} x_{k_j,j} + \left(1 - \frac{1}{\theta_{k'}}\right) \frac{1}{\theta_{k'-1}\theta_{k'-2}} \sum_{k''=k_j+2}^{k'} \theta_{k''-2} z_{k''-1} + \frac{1}{\theta_{k'}} z_{k',j} \\ &= \frac{\theta_{k_j}\theta_{k_j-1}}{\theta_{k'}\theta_{k'-1}} x_{k_j,j} + \frac{1}{\theta_{k'}\theta_{k'-1}} \sum_{k''=k_j+2}^{k'} \theta_{k''-2} z_{k''-1} + \frac{1}{\theta_{k'}} z_{k',j} \\ &= \frac{\theta_{k_j}\theta_{k_j-1}}{\theta_{k'}\theta_{k'-1}} x_{k_j,j} + \frac{1}{\theta_{k'}\theta_{k'-1}} \sum_{k''=k_j+2}^{k'+1} \theta_{k''-2} z_{k''-1}. \end{aligned}$$

The first equality is due to the definition of $x_{k'}$. The second equality follows from the assumption of induction. The third equality holds by Lemma A.1. This shows that (9) holds.

Next we show that

$$z_{k'-1,j} = \frac{1}{1+\eta\theta_{k'-1}\theta_{k'-2}\lambda_2}\mathrm{soft}\left(z_{0,j} - \eta g^{\mathrm{sum}}_{k_j,j} - \eta(\theta_{k'-1}\theta_{k'-2} - \theta_{k_j}\theta_{k_j-1})\widetilde{\nabla}_j, \eta\theta_{k'-1}\theta_{k'-2}\lambda_1\right), \tag{10}$$

for $k' \in [k_j+2, k]$.

By the definition of $z_{k'-1}$, we have that

$$\begin{aligned} z_{k'-1,j} &= \mathrm{prox}_{\eta\theta_{k'-1}\theta_{k'-2}R}(z_0 - \eta\theta_{k'-1}\theta_{k'-2}\bar{g}_{k'-1})_j \\ &= \frac{1}{1+\eta\theta_{k'-1}\theta_{k'-2}\lambda_2}\mathrm{soft}(z_{0,j} - \eta\theta_{k'-1}\theta_{k'-2}\bar{g}_{k'-1,j}, \eta\theta_{k'-1}\theta_{k'-2}\lambda_1) \end{aligned}$$

From Lemma A.4, we see that

$$\begin{aligned} \theta_{k'-1}\theta_{k'-2}\bar{g}_{k'-1,j} &= \sum_{k''=1}^{k'-1}\theta_{k''-1}g_{k'',j} \\ &= \sum_{k''=1}^{k_j}\theta_{k''-1}g_{k'',j} + \left(\sum_{k''=k_j+1}^{k'-1}\theta_{k''-1}\right)\widetilde{\nabla}_j \\ &= g^{\mathrm{sum}}_{k_j,j} + (\theta_{k'-1}\theta_{k'-2} - \theta_{k_j}\theta_{k_j-1})\widetilde{\nabla}_j. \end{aligned}$$

The first and third equality are due to Lemma A.2. The second equality holds because $g_{k''-1,j} = \widetilde{\nabla}_j$ for $k'' \in [k_j+1, k-1]$ by the assumption. This shows that (10) holds. Observe that

$$\begin{aligned} z_{k'-1,j} &= \frac{1}{1+\eta\theta_{k'-1}\theta_{k'-2}\lambda_2}\mathrm{soft}\left(z_{0,j} - \eta g^{\mathrm{sum}}_{k_j,j} - \eta(\theta_{k'-1}\theta_{k'-2} - \theta_{k_j}\theta_{k_j-1})\widetilde{\nabla}_j, \eta\theta_{k'-1}\theta_{k'-2}\lambda_1\right) \\ &= \frac{1}{1+\eta\theta_{k'-1}\theta_{k'-2}\lambda_2}\mathrm{sign}\left(z_{0,j} - \eta g^{\mathrm{sum}}_{k_j,j} - \eta(\theta_{k'-1}\theta_{k'-2} - \theta_{k_j}\theta_{k_j-1})\widetilde{\nabla}_j\right) \\ &\quad \times \max\left\{\left|z_{0,j} - \eta g^{\mathrm{sum}}_{k_j,j} - \eta(\theta_{k'-1}\theta_{k'-2} - \theta_{k_j}\theta_{k_j-1})\widetilde{\nabla}_j\right| - \eta\theta_{k'-1}\theta_{k'-2}\lambda_1, 0\right\} \\ &= \begin{cases} \frac{1}{1+\eta\theta_{k'-1}\theta_{k'-2}\lambda_2}(z_{0,j} - M^+_{k',j}) & (z_{0,j} > M^+_{k',j}) \\ 0 & (M^-_{k',j} \le z_{0,j} \le M^+_{k',j}) \\ \frac{1}{1+\eta\theta_{k'-1}\theta_{k'-2}\lambda_2}(z_{0,j} - M^-_{k',j}) & (z_{0,j} < M^-_{k',j}) \end{cases}, \end{aligned}$$

where $M^{\pm}_{k',j} = \eta\theta_{k'-1}\theta_{k'-2}(\widetilde{\nabla}_j \pm \lambda_1) + \eta g^{\mathrm{sum}}_{k_j,j} - \eta\theta_{k_j}\theta_{k_j-1}\widetilde{\nabla}_j$. We define the real valued functions $M^{\pm}$ as follows:

$$M^{\pm}_j(x) \stackrel{\mathrm{def}}{=} (c_1 \pm c_2)x^2 - (c_1 \pm c_2)x + c_3,$$

where $c_1 = \frac{\eta\widetilde{\nabla}_j}{4}$, $c_2 = \frac{\eta\lambda_1}{4}$ and $c_3 = \eta g^{\mathrm{sum}}_{k_j,j} - \eta\theta_{k_j}\theta_{k_j-1}\widetilde{\nabla}_j$ Then we see that $M^{\pm}_j(k') = M^{\pm}_{k',j}$. Let

$$\begin{aligned} D_{\pm} &\stackrel{\mathrm{def}}{=} (c_1 \pm c_2)^2 + 4(c_1 \pm c_2)(z_{0,j} - c_3), \\ s^+_{\pm} &\stackrel{\mathrm{def}}{=} \frac{c_1 + c_2 \pm \sqrt{D_+}}{c_1 + c_2}, \\ s^-_{\pm} &\stackrel{\mathrm{def}}{=} \frac{c_1 - c_2 \pm \sqrt{D_-}}{c_1 - c_2}, \end{aligned}$$

where if $s_{\pm}^{\pm}$ are not well defined, we simply assign 0 (or any number) to $s_{\pm}^{\pm}$. We can easily show that the following results:
1) If $c_1 > c_2$, then

$$z_{0,j} > M_j^+(x) \iff \begin{cases} x \in \emptyset & (D_+ \le 0) \\ s_-^+ < x < s_+^+ & (D_+ > 0) \end{cases},$$
$$z_{0,j} < M_j^-(x) \iff \begin{cases} x \in \mathbb{R} & (D_- \le 0) \\ x < s_-^- \vee x > s_+^- & (D_- > 0) \end{cases}.$$

2) If $c_1 = c_2$, then

$$z_{0,j} > M_j^+(x) \iff \begin{cases} x \in \emptyset & (c_2 = 0 \wedge z_{0,j} \le c_3) \\ x \in \mathbb{R} & (c_2 = 0 \wedge z_{0,j} > c_3) \\ x \in \emptyset & (c_2 > 0 \wedge D_+ \le 0) \\ s_-^+ < x < s_+^+ & (c_2 > 0 \wedge D_+ > 0) \end{cases},$$
$$z_{0,j} < M_j^-(x) \iff \begin{cases} x \in \mathbb{R} & (z_{0,j} < c_3) \\ x \in \emptyset & (z_{0,j} \ge c_3) \end{cases}.$$

3) If $|c_1| < c_2$, then

$$z_{0,j} > M_j^+(x) \iff \begin{cases} x \in \emptyset & (D_+ \le 0) \\ s_-^+ < x < s_+^+ & (D_+ > 0) \end{cases},$$
$$z_{0,j} < M_j^-(x) \iff \begin{cases} x \in \emptyset & (D_- \le 0) \\ s_-^- < x < s_+^- & (D_- > 0) \end{cases}.$$

4) If $c_1 = -c_2$, then

$$z_{0,j} > M_j^+(x) \iff \begin{cases} x \in \emptyset & (z_{0,j} \le c_3) \\ x \in \mathbb{R} & (z_{0,j} > c_3) \end{cases}$$
$$z_{0,j} < M_j^-(x) \iff \begin{cases} x \in \mathbb{R} & (c_2 = 0 \wedge z_{0,j} < c_3) \\ x \in \emptyset & (c_2 = 0 \wedge z_{0,j} \ge c_3) \\ x \in \emptyset & (c_2 > 0 \wedge D_- \le 0) \\ s_-^- < x < s_+^- & (c_2 > 0 \wedge D_- > 0) \end{cases}.$$

5) If $c_1 < -c_2$, then

$$z_{0,j} > M_j^+(x) \iff \begin{cases} x \in \mathbb{R} & (D_+ \le 0) \\ x < s_-^+ \vee x > s_+^+ & (D_+ > 0) \end{cases},$$
$$z_{0,j} < M_j^-(x) \iff \begin{cases} x \in \emptyset & (D_- \le 0) \\ s_-^- < x < s_+^- & (D_- > 0) \end{cases}.$$

The lazy update rules of $x_{k-1,j}$ is derived by combining (9) with these results and noting that $k' \in [k_j + 2, k]$. Finally, combining the definition $y_{k,j} = (1 - 1/\theta_k)x_{k-1,j} + (1/\theta_k)z_{k-1,j}$ with (10) gives the lazy update of $y_{k,j}$. The update rule of $z_{k,j}$ is obvious from the proof of (10). □

C. Experimental Details

In this section, we give the experimental details and also comment on the adaptivity of SVRG to local strong convexity.

The details of the implemented algorithms and their parameter tunings were as follows:

For non-strongly convex cases ($(\lambda_1, \lambda_2) = (10^{-4}, 0)$),

- SVRG^{++} (Allen-Zhu & Yuan, 2016) with default initial epoch length $m = n/(4b)$ [7]. We tuned only the learning rate.
- AccProxSVRG (Nitanda, 2014). We tuned the epoch length, the constant momentum rate and the learning rate, and additional dummy ℓ_2 regularizer weight for handling a non-strongly convex objective.
- UC (Lin et al., 2015) + SVRG (Xiao & Zhang, 2014) with default epoch length $m = 2n/b$ [8]. We tuned κ in (Lin et al., 2015) and the learning rate. We fixed $\eta = 1$ in the algorithm of UC (note that η is not learning rate).
- UC + AccProxSVRG. We tuned κ in (Lin et al., 2015), the epoch length, the constant momentum rate and the learning rate. We fixed $\eta = 1$ in the algorithm of UC (note that η is not learning rate).
- APCG (Lin et al., 2014a). We tuned the convexity parameter of the dual objective and the learning rate, and additional dummy ℓ_2 regularizer weight for handling a non-strongly convex objective.
- Katyusha$^{\mathrm{ns}}$ (Allen-Zhu, 2016) with default epoch length $m = 2n/b$ and Katyusha momentum $\tau_2 = 1/2$ following the suggestion of (Allen-Zhu, 2016). We tuned only the learning rate. We did not adopt AdaptReg scheme because Katyusha with AdaptReg was always a bit slower than vanilla Katyusha in our experiments.
- DASVRDA$^{\mathrm{ns}}$ with epoch length $m = n/b$ and $\gamma = \gamma_*$. We tuned only the learning rate.
- Adaptive Restart DASVRDA with epoch length $m = n/b$ and $\gamma = \gamma_*$. We tuned only the learning rate. We used the gradient scheme for the adaptive restarting, that is we restart DASVRDA$^{\mathrm{ns}}$ if $(\widetilde{y}_s - \widetilde{x}_s)^\top (\widetilde{y}_{s+1} - \widetilde{x}_s) > 0$.

For strongly convex cases ($(\lambda_1, \lambda_2) = (10^{-4}, 10^{-6}), (0, 10^{-6})$),

- SVRG (Xiao & Zhang, 2014) with default epoch length $m = 2n/b$. We tuned only the learning rate.
- AccProxSVRG (Nitanda, 2014). We tuned the epoch length, the constant momentum rate and the learning rate.
- UC (Lin et al., 2015) + SVRG (Xiao & Zhang, 2014) with default epoch length $m = 2n/b$ [4]. We tuned κ, q in (Lin et al., 2015) and the learning rate.
- UC + AccProxSVRG. We tuned κ, q in (Lin et al., 2015), the epoch length, the constant momentum rate and the learning rate.
- APCG (Lin et al., 2014a). We tuned the convexity parameter of the dual objective and the learning rate.
- Katyusha (Allen-Zhu, 2016) with default epoch length $m = 2n/b$ and Katyusha momentum $\tau_2 = 1/2$ following the suggestion of (Allen-Zhu, 2016). We tuned τ_1 in (Allen-Zhu, 2016) and the learning rate.
- DASVRDA$^{\mathrm{sc}}$ with epoch length $m = n/b$ and $\gamma = \gamma_*$. We tuned the fixed restart interval S and the learning rate.
- Adaptive Restart DASVRDA with epoch length $m = n/b$ and $\gamma = \gamma_*$. We tuned only the learning rate. We use the gradient scheme for the adaptive restarting, that is we restart DASVRDA$^{\mathrm{ns}}$ if $(\widetilde{y}_s - \widetilde{x}_s)^\top (\widetilde{y}_{s+1} - \widetilde{x}_s) > 0$.

For tuning the parameters, we chose the values that led to the minimum objective value. We selected the learning rates from the set $\{10^p, 2 \times 10^p, 5 \times 10^p \mid p \in \{0, \pm 1, \pm 2\}\}$ for each algorithm. We selected the epoch lengths from the set $\{n \times 10^{-k}, 2n \times 10^{-k}, 5n \times 10^{-k} \mid k \in \{0, 1, 2, 3\}\}$ and the momentum rates from the set $\{1 - 10^{-k} \mid k \in \{1, 2, 3, 4\}\}$ for AccProxSVRG. We chose the additional dummy ℓ_2 regularizer weights from the set $\{10^{-k}, 0 \mid k \in \{4, 5, 6, 8, 12\}\}$

[7] In (Allen-Zhu & Yuan, 2016), the authors have suggested a default initial epoch length $m = n/4$. Since we used mini-batches with size b in our experiments, it was natural to use $m = n/(4b)$. We made sure that using this epoch length improved the performances in all settings.

[8] In (Xiao & Zhang, 2014), the authors has suggested a default initial epoch length $m = 2n$. Since we used mini-batches with size b in our experiments, it was natural to use $m = 2n/b$. We made sure that using this epoch length improved the performances in all settings.

for AccSVRG and APCG. We selected κ, q from the set $\{10^{-k} \mid k \in \{1,2,3,4,5,6\}\}$ for UC. We chose the convexity parameter from the set $\{10^{-k} \mid k \in \{3,4,5,6,7\}\}$ for APCG. We selected τ_1 from the set $\{10^{-k}, 2\times 10^{-k}, 5\times 10^{-k} \mid k \in \{1,2,3\}\}$ for Katyusha. We selected the restart interval from the set $\{10^k, 2\times 10^k, 5\times 10^k \mid k \in \{0,1,2\}\}$ for DASVRDA$^{\mathrm{sc}}$.

We fixed the initial points $0 \in \mathbb{R}^d$ for all algorithms.

For a fair comparison, we used uniform sampling for all algorithms, because AccProxSVRG does not support non-uniform sampling.

Remark on the adaptivity of SVRG to local strong convexity

For an $L1$-reguralization function, the finite sum is often locally strongly convex (or restricted strongly convex) over the set of sparse solutions[9](see (Agarwal et al., 2010; 2012; Lin et al., 2014b)). For example, the logistic models in our experiments satisfy this property. In our elastic net setting, the regularization parameter of $L1$-regularalization (10^{-4}) is much larger than the one of $L2$-reguralization (10^{-6}). Thus, exploiting the local strong convexity is important for faster convergence. In (Allen-Zhu & Yuan, 2015), the author has proposed a modified SVRG algorithm for strongly convex objectives. The essential difference from vanilla SVRG is only using a weighted average of inner updated solutions for the output solution, rather than using uniform average. Their analysis assumes the strong convexity of the finite sum rather than the regularizer and it is possible that the algorithm exploits the local strong convexities of loss functions. More importantly, their algorithm only uses the strong convexity parameter in the weights for the output solution. Specifically, the output solution has the form $\widetilde{x} = (1/\sum_{k=1}^m (1-\mu\eta)^{-k}) \sum_{k=1}^m (1-\mu\eta)^{-k} x_k$, where $\eta = O(1/L)$. Thus if the epoch length m is relatively small, the output solution is almost same as the uniformly averaged one (Allen-Zhu & Yuan (2015) assumes $m = O(L/\mu)$ but we have discovered that using $m = O(n/b)$ still guarantees that the algorithm achieves the convergence rate $O((n + bL/\mu)\log(1/\varepsilon))$, that is same as the one of vanilla SVRG in mini-batch settings). The reason why using $m = O(n/b)$ rather than $m = O(n)$ gave faster convergence in our experiments is probably because using $m = O(n/b)$ ensured that the weighted average was nearer to the uniform average than using $m = O(n)$. Therefore, we can say that vanilla SVRG with epoch length $m = O(n/b)$ is almost adaptive to local strong convexity. In contrast, Katyusha uses the strong convexity parameter in the momentum rate and is quite sensitive. Thus it seems to be difficult for Katyusha to exploit local strong convexity.

[9]We say that function F is locally strongly convex with respect to $L1$-regularizer if there exist $\gamma > 0$ and $\tau > 0$ such that $F(x) - F(y) - \langle \nabla F(y), x - y\rangle \geq (\gamma/2)\|x-y\|_2^2 - \tau\|x-y\|_1^2$ for any $x, y \in \mathbb{R}^d$. Thus, if F is locally strongly convex, then for any $y \in \mathbb{R}^d$ and $\gamma' < \gamma$, F is $\gamma - \gamma'$-strongly convex over the set $\{x \in \mathbb{R}^d \mid \|x-y\|_1 \leq \sqrt{\gamma'/(2\tau)}\|x-y\|_2\}$. This intuitively means that if F is locally strongly convex, F is strongly convex over the set of sparse solutions whose supports do not fluctuate from each other.